

GET THE MOST OUT OF YOUR NEW AMD RYZEN

apoc

Nov 2019
Issue 473

YOUR EXPERT

GUIDE TO TODAY'S TECH



OVERCLOCK & TWEAKING GUIDE



SECURE

TIRED OF LOSING LOGINS?
WE SHOW YOU THE SECRET TO BETTER SECURITY

[PROCESSING] 02
HACK ATTEMPT
SCAN NETWORK
USER SAFE
FAILED

BEST RAM UPGRADES


FACEBOOK LIBRE
IS THE WORLD READY FOR ZUCK BUCKS?

WINDOWS STORAGE SPACES
FULL GUIDE TO MICROSOFT'S LITTLE-KNOWN WONDER



BEST RAM UPGRADES

- * SPEED VS. CAPACITY
- * UNLEASH YOUR CPU & GPU
- * HOW MUCH RAM DO YOU REALLY NEED?



TWIN-SCREENS:
INNOVATION OR EXCESS?

TESTED:
ASUS ZENBOOK DUO
PLUS, DELL XPS 15, GOOGLE NEST HUB, SAMSUNG SPACE MONITOR, ACER PREDATOR HELIOS 300



GOOGLE NEST HUB
PAGE 40

FUTURE

AUS\$99.95 INC GST
NZ\$10.90 INC GST
PRINT POST
APPROVED: 100000179
ISSN: 0728-4415

11

313006 001362

NITRO

NITRO 5



acer



Powered by Intel® Core™ Processors.
Intel Inside®. Extraordinary
Performance Outside.



NITRO 5

Combined power and
performance of a 9th Gen
Intel® Core™ i7 processor
and up to NVIDIA®
GeForce® GTX 1660 Ti



COOLNESS UNDER FIRE

Chill out with twin
fans, Acer CoolBoost™
technology and a dual
exhaust port design



UP TO
32GB
RAM

UP TO
1TB
HDD

UP TO
512GB
SSD

SUPERCHARGED CONNECTIVITY AND STREAMING

Control your network with
Killer™ Ethernet E25002
and effortlessly stream with
Gigabit Wi-Fi

VISUAL INTENSITY

Sharp, vibrant images
of a 15.6" FHD IPS
display with up to
144hz refresh rates

DOMINATING YOUR GAME

Available now at

**Harvey
Norman**

**JB
HI-FI**

**PLE
COMPUTERS**

**mwave
com.au**

CENTRE COM
All Centre Zone and Computer & IT Stores

**PC
CC**

Umart

scorptec computers



apc

Future Publishing Australia, PO Box 1077 Mount Street,
North Sydney, NSW 2059
Subscription enquiries: Please call CRM: (02) 8227 6486
Tel: (02) 9955 2677 Fax: (02) 9955 2688
Email: apcmag@futurenet.com
Web: www.apcmag.com

A nice time

To be playing with tech.



Choosing which products to review in APC isn't always straightforward. Sometimes it's dictated by the seasons – there are quiet times, and there are months

when every company seems to have something new. Then you get the really crazy periods when a flood of products are unleashed because there's a new version of the underlying technology that is used by several manufacturers. We see this in motherboards, for example, when a new chipset comes out, along come dozens of new boards. That's when we do group tests instead of stand-alone reviews.

This is a good month for reviews. I was particularly excited to get the new Asus ZenBook Duo in. I first saw this innovative dual-screen laptop at Computex earlier this year and I, along with the rest of the APC team in attendance, all agreed it was the most impressive single product of the entire show. Well, we've had one in the Labs for a couple of weeks and Joel has been digging deep. What does he think? Turn to page 26 for the review!

We've also got the new Dell XPS 15 in (page 28). The entire XPS series from Dell has been long admired by many since it almost single-handedly brought sexy back to PC

laptops. I do believe that Dell's uncompromising desire to create a slim, slick and stylish laptop range has driven many other companies to do the same; after all, the XPS range has always sold extremely well and with Dell establishing a market for such machines, others were bound to follow. The latest model is in, and perhaps the review can tempt you to consider one, or even better – your company's IT department may want to treat you to one of the nicest PCs available today.

Google's Nest Hub is here (page 40). As we have observed, the market for Google Home display devices hasn't taken off – at least with partner manufacturers. Only JBL, LG and Lenovo have released such devices, and perhaps, we wonder, Google grew impatient with the lack of energy in this segment so put this product together. Which is not to suggest it was a rush job; as you'll read it's an outstanding product. That belongs in the kitchen.

Mostly, though, my heart was warmed to receive Microsoft's new Intellimouse. It's just like the original from the late 1990s, but with new switches and a new sensor. It's a classic reborn and I love it. RGB? Look elsewhere. See you next month!

BEN MANSILL

EDITOR

ben.mansill@futurenet.com



Editorial

Editor: Ben Mansill
Senior Journalist: Shaun Prescott
Journalist: Joel Burgess
Journalist: Harry Domanski
Journalist: Stephen Lambrechts
Journalist: Sharmishta Sarkar
Creative Director: Troy Coleman
Designer: Sharnie Swinnerton

Contributors

Anthony Agius, George Cairns, Natalie Clayton, Alex Cox, Dr Ian Cutress, James Davenport, Sean Endicott, Christian Guyton, Lindsey Handmer, Kenny Hemphill, Marshall Honorof, Anthony Howes, Phil Iwaniuk, Hollin Jones, Rick Lane, Stephen Lambrechts, Chris Lloyd, Neil Mohr, James O'Connor, Nick Peers, Jeremy Peel, Mike Prospero, Les Pounder, Chris Szewczyk, Mayank Sharma, Jarren Walton, Darren Yates, Robert Zak.

Photography

All copyrights and trademarks are recognised and respected

Advertising

Group Advertising Manager: Cameron Ferris
cameron.ferris@futurenet.com

Management

Managing Director: Neville Daniels
Sales Director: Paul Marttila

Printed by Blue Star PRINT (NSW)

Distributed in Australia and NZ by Gordon and Gotch
www.gordongotch.com.au

ISSN 0728-441

We are committed to only using magazine paper which is derived from responsibly managed, certified forestry and chlorine-free manufacture. The paper in this magazine was sourced and produced from sustainable managed forests, conforming to strict environmental and socioeconomic standards. The manufacturing paper mill holds full FSC (Forest Stewardship Council) or PEFC certification and accreditation.

All contents © 2018 Future Publishing Australia or published under licence. All rights reserved. No part of this magazine may be used, stored, transmitted or reproduced in any way without the prior written permission of the publisher. Future Publishing Limited (company number 2008885) is registered in England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All information contained in this publication is for information only and is, as far as we are aware, correct at the time of going to press. Future cannot accept any responsibility for errors or inaccuracies in such information. You are advised to contact manufacturers and retailers directly with regard to the price of products/ services referred to in this publication. Apps and websites mentioned in this publication are not under our control. We are not responsible for their contents or any other changes or updates to them. This magazine is fully independent and not affiliated in any way with the companies mentioned herein.

If you submit material to us, you warrant that you own the material and/or have the necessary rights/permissions to supply the material and you automatically grant Future and its licensees a licence to publish your submission in whole or in part in any/all issues and/or editions of publications, in any format published worldwide and on associated websites, social media channels and associated products. Any material you submit is sent at your own risk and, although every care is taken, neither Future nor its employees, agents, subcontractors or licensees shall be liable for loss or damage. We assume all unsolicited material is for publication unless otherwise stated, and reserve the right to edit, amend, adapt all submissions.

Privacy statement

If you provide information about yourself this will be used to provide you with products or services you have requested. We may supply your information to contractors to enable us to do this. Future Publishing Australia will also use your information to inform you of other publications, products, services and events. Future Publishing Australia may also give your information to organisations that are providing special prizes or offers and are clearly associated with the Reader Offer. Unless you tell us not to, Future Publishing Australia may give your information to other organisations that may use it to inform you of other products, services or events. If you would like to gain access to the information Future Publishing Australia holds about you, please contact us.




Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief executive Zillah Byng-Thorne
Non-executive chairman Peter Allen
Chief financial officer Penny Larkin-Brand

Tel +44 (0)1225 442 242

BLANK CANVAS,
UNLIMITED POSSIBILITIES



H200 BLACK & SNOW
RGB TEMPERED GLASS CASE

exclusive software

» HOW TO DOWNLOAD YOUR FREE FULL-VERSION PROGRAMS

ISSUE 473, NOVEMBER 2019

Please note that these exclusive downloads will only be available for a limited time, from **07/10/2019 to 24/11/2019**




FLUSH YOUR REGISTRY WELL AND GOOD

OPTIMISE YOUR PC'S PERFORMANCE.

FULL VERSION WORTH \$20

The Windows registry is, in a way, the administrative center of your system. Different settings and configuration data for Windows and most of your program's settings, like for example autostart settings, are saved there. Over time errors creep in the registry. One result is that your PC will become slower, e.g. when starting the system. Problems when uninstalling software and a large number of error messages could point to problems in the Windows registry. With the Ashampoo Registry Cleaner you will easily and safely delete redundant or damaged registry entries. Faulty links in the registry are also repaired. The real-time backup in the Ashampoo Registry Cleaner offers a maximum of safety and enables the restoring of the windows registry into its original status anytime.

DOWNLOAD LINK: www.ashampoo.com/lpa/apcmagazine



SPARE YOURSELF FROM INTERNET VILLAINS


ELIMINATE AGGRESSIVE VIRUSES WITH ABELSOFT ANTIRANSOMWARE.

FULL VERSION WORTH \$26

With one click, Data2Pocket searches your entire PC for pictures, videos and audio files, then drags it all to your selected USB stick or, if necessary, to a larger USB disk. You need no longer worry that you might forget an important photo or song: Data2Pocket finds them all in no time. Never again will you have to spend hours searching your computer for your most personal files: Simply insert the stick, start the program and then put the stick into your trouser pocket.

DOWNLOAD LINK: www.apcmag.com/exclusives

DOWNLOAD ZIP PASSWORD: **September19i** (note this is case sensitive)



DISCLAIMER Future Publishing is not the licensor of the exclusive software or any documentation included with it. Rather, Future is merely a distributor of the software, and your use of the software is subject to any accompanying third-party licence terms. You must carefully read and comply with any such third-party licence terms, together with all instructions and README files that come with the software. Accordingly, to the maximum extent permitted by law, all software is provided by Future 'as is' and without warranty, either express or implied, and Future will not be liable for any damage that you may incur as a result of using any software downloads. You should back up any important system and data files before using any downloaded software. We recommend that you do not use the exclusive software on a production machine. Nothing in this disclaimer excludes any warranty that may be implied by statute, which may include the Trade Practices Act. While we have taken all reasonable steps to check the software downloads for viruses, we cannot guarantee that it is free from viruses or other harmful code and you should check each download using a virus scanner complete with the latest antivirus updates before use.


PERFECT CHOICE, GENUINE VALUE

5GHz
1800Mbps

2.4GHz
600Mbps

**BiPAC 8900AX-2400**

Wireless-AC 2400Mbps LTE
VDSL2/ADSL2+ VPN Firewall Router

au.billion.com

- ▶ VDSL(FTTC)/ADSL/EWAN for all connections
- ▶ Turbo-charged wireless AC with future-proof MU-MIMO support
- ▶ 4X4 (5 GHz) 3x3 (2.4GHz)
- ▶ Powerful 2.4GHz with turboQAM for best speed & coverage

Contents

56

SECURE

TIRED OF LOSING LOGINS?
WE SHOW YOU THE SECRET
TO BETTER SECURITY

46

LONG TERM MEMORY

Desktop memory is becoming more affordable; Christian Guyton examines what that could mean for you.



■ features

46 MEMORY BUYER'S GUIDE

Prices are falling and the new breed of CPUs have very particular requirements

56 PERFECT PASSWORD PRIMER

It's time to get your passwords sorted so things are simpler, easier to manage and *much* more secure

68 MICROSOFT STORAGE SPACE

At the edge of the final storage frontier, Anthony Howes looks at how the universe's greatest storage system is about to unfold.

00 FACEBOOK'S NEW MONEY

Chris Lloyd examines the social media giant's impact on the finance industry with Libra

■ technotes

12 NEWS

The latest developments in the tech world

16 TECH BRIEF

Chip yields, binning, and the silicon lottery

18 GADGETS

Lovely little bits and pieces

20 TWO BITS

Are high end desktop platforms still relevant?

21 END USER

Online anger may be silly, but the Epic Games Store is undeniably bad

22 RANDOM ACCESS

Is OLED on laptops shaping up as a winning tech?

23 NUMBER CRUNCH

We look at the numbers driving the big tech news

"Human laziness, complacency, general failings, and unending ingenuity to subvert systems have undermined passwords forever. *page 56*



40 GOOGLE NEST HUB



28 DELL XPS 15



35 MICROSOFT INTELLIMOUSE



32 SYNOLOGY DS620SLIM



34 TP LINK ARCHER MR600 4G+ ROUTER

the lab

LATEST REVIEWS

- 26** Asus ZenBook Duo UX581 laptop
- 28** Dell XPS 15
- 29** Acer Predator Helios 300
- 30** Asus Zenbook 14 UX433F
- 31** Samsung Space Monito
- 32** Synology DS620slim
- 33** Draytek VigorAP903
- 34** TP Link Archer MR600 4G+
- 34** Thermaltake ToughRAM
- 35** Microsoft Intellimouse
- 35** Logitech G502 Lightspeed
- 36** Steel Series Apex 7
- 37** Thermaltake TT Floe RGB 360 AIO
- 38** Logitech 915 keyboard
- 39** Motorola G7 and G7 Plus
- 40** Google Nest Hub

PC BUILDER

- 00** System News
- 00** Market Watch
- 00** Blueprint
- 00** PC Builder Masterclass

how to

QUICK TIPS

- 82** We fix your tech problems

TUTORIALS

- 86** Get the most from your AMD CPU
- 88** Create a table of contents in Pages
- 90** Enhance sound in iMovie
- 92** Store and sync notes across your PC and devices

MASTERCLASSES

- 96** Get started with Augmented Reality on Android
- 100** Create a Pi OLED monitor
- 102** How and where you use machine-learning

downtime

104 GAMING NEWS

What's happening in the world of fun

106 GAMING REVIEWS

High-performance playtime

108 GAME CHANGER

Enter the Matrix

114 CHIP CHAT

Quirky news from the world of geekdom

technotes

» INSIDE APC



theLab » latest reviews

"If you're happy with 1080p gaming the Nvidia GTX 1050 GPU is serviceable netting around 30fps for current games on Ultra 1080p settings."

LABS FROM \$199 | WWW.DELL.COM

Dell XPS 15 OLED

Dell goes after the detail-oriented with its latest XPS 15 OLED professional ultrabook.

The new Dell XPS 15 can be configured with an overclockable Core i9-9900K CPU, 32GB of RAM and a 1TB NVMe SSD in RAID 0 configuration for \$5,169. This is a powerful CPU that we really only seen on overpowered gaming rigs and while it's great to have this chip on Dell's ultrabook it requires a generous cooling infrastructure to work well and we're not fully convinced that Dell has nailed the thermal management just yet.

28 www.apcmag.com

Full Adobe RGB would make it the broadest colour gamut of the laptop OLED screens we've seen to date, however you'll have to take Dell's word for it since it doesn't offer any third-party calibration or HDR certifications. Notebookcheck measured the screen's capability at 81% Adobe RGB 1995, which is a bit of a claim but still much better than the 68% scored by the Gigabyte Aero 15 OLED. While this device does offer Dolby Vision HDR movie playback, the screen output seems to be locked to the Intel UHD Graphics 630, so you won't be able to game in HDR, even though the screen and GPU are both capable of it. It's not uncommon for laptop vendors to configure GPUs in this way in order to minimise battery expenditure when they're not in use, but it does isolate a

number of GPU functionalities including HDR gaming from the screen. We tested the Dell XPS 15 with an Intel Core i7-9750H CPU and 32GB RAM allocation that is listed on Dell's website at \$4,209, but we've seen on sale for \$3,654. While raw CPU and media encoding scores of the XPS 15 were between 10 and 17% better than our average for this CPU, the general work tasks were between 3 and 20% behind other comparable devices. This is potentially due to some throttling over longer CPU-heavy tasks since we were regularly able to get the CPU to spike to 100°C. The i9 model, presumably, uses a similar thermal infrastructure to the XPS 15, which makes us a little worried about the potential for the i9 to offer any performance benefits at all. If you're happy with 1080p gaming the Nvidia

GTX 1050 GPU is serviceable, netting around 30fps for current games on Ultra 1080p settings. If you're playing slightly older games or you're happy to compromise on graphical settings then you can even get some titles running at 30fps in 4K. Despite a large 57Wh battery, the device only lasts four hours and seven minutes on the PCMark 8 battery benchmark and less for movie playback, which means it'll be a bit shy of the full day mark if you use it sparingly. **Joel Burgess**

Verdict

A great option for professionals that you use it sparingly. **★★★★★**

Inside APC

Find out all about APC's editorial policies, test practices, how to read the benchmark results, and more.

APC is Australia's oldest consumer technology magazine – having been consistently in print for over 35 years, since our first issue way back in May 1980 – and we take that heritage and responsibility very seriously. While our focus is obviously on the personal computer – it's in our name, after all – the very definition of the PC has changed and shifted markedly since the early 1980s. As such, we touch on many other areas of tech, too, from smartphones and apps to peripherals, accessories, online services and beyond. We have two main goals: to track down the best of modern tech and also to help our readers make the most of it.

We're also an open church in terms of platforms. We know most people aren't wed to a single brand's products and use a variety of devices. And, like you, APC's journalists want to know what's good in tech – no matter what platform it resides on.

INDEPENDENT REVIEWS

Championing technology doesn't mean we're unrelenting yes-men and -women, however, and APC aims to be as objective as possible in all our coverage. That means identifying the best products from multiple perspectives – the best performance, best value and best features and, ideally, the products that offer the best mix of these three.

As a matter of policy, reviews published in APC are not shared with product-makers prior to print. We will contact vendors under certain conditions; for example, if we have a problem testing a product that seems to indicate it may be faulty, or to invite a vendor to clarify how a particular feature works. If an APC reviewer has any potential conflicts of interest involving a brand, the review will always be assigned to another writer.

LABS TESTING

Despite being a small magazine with limited resources, APC still strives to conduct the most rigorous, objective scientific tests and benchmarks we can so as to make our reviews as unbiased as possible. We use a variety of tools and programs for this, including many freely available benchmark suites for assessing media encoding, general system performance, gaming and battery life.

In most cases, for the benchmark results published in APC, you can assume that higher is better. There are certain tests that deviate from this rule where the opposite is true; in those cases, we've flagged the results with a note explaining as such.

We use both tables and graphs for displaying results; the latter are our preference due to their ease-of-readability, but tables are more compact, so we use these in cases where thoroughness is preferred.



ELEGANCE MAKES GLORY

RADEON RX 5700 XT TAICHI X 8G OC SERIES


ASRock

www.asrock.com

POLYCHROME
SYN C

AMD
RADEON RX
5700 XT

The brand and product names are trademarks of their respective companies.



Next-gen chips could be made from carbon nanotubes


GIVING SILICON THE FLICK.

Researchers at the University of Illinois are working on a computer chip that uses carbon nanotubes instead of silicon as material for the transistors. This would allow processors to run up to three times faster yet use less than a third the power of silicon chips, which are nearing the physical limits of achievable performance.

AMD responds to customers complaining Ryzen doesn't reach max boost speeds

ADMITS THERE WAS AN ISSUE AND DELIVERS BIOS FIX.

Following complaints from users that their new Ryzen 3000 CPUs weren't reaching advertised boost speeds, AMD has responded with a BIOS patch that it has released to its motherboard partners. The fix targets the fastest cores with lightly-threaded applications, allowing a more balanced workload across all cores.



Kaspersky testing an eSports cheat detector

COULD BE USED IN TOURNAMENTS.

The company announced a beta version of Kaspersky Anti-Cheat, a tool specifically made for eSports organisations to ensure fair play. Its cloud-based solution is installed in players' systems, then "game process information is collected, sent to the Kaspersky Anti-Cheat cloud and analysed for suspicious events." A referee then receives reports in real time through a web interface, leaving the ref and organisation to decide what actions to take.

nbnTM



NBN Co has "no plans" for a Netflix tax – proposes new 1Gbps plan

REST EASY, STREAMERS.

In a past consultation paper, NBN Co floated the idea of raising the prices for Australian internet users based on the type of data they were using rather than the amount. Essentially, this translated to what was dubbed the 'Netflix tax', as it was predominantly targeting streaming media during peak hour. The proposal was met with wide backlash and considered an attack on net neutrality, but NBN backedpedalled on the seriousness of the proposal. Now, NBN Co's chief customer officer Brad Whitcomb has said there are "no plans" to charge users based on their streaming, with the most recent consultation paper revealing that only two internet service providers (ISPs) actually supported charging users based on their usage, but they proposed video streaming be charged less. In order to accommodate these data-hungry users, NBN co has proposed new 250Mbps and 1Gbps speed tiers, although it's ultimately up to the ISPs whether or not these speeds will be passed on to consumers.

"In order to accommodate these data-hungry users, NBN co has proposed new 250Mbps and 1Gbps speed tiers."

Telstra finally joined the eSIM club with the iPhone 11 launch


SWITCHING TELCOS IS NOW EVEN ESIMPLER.

Telstra rolled out eSIM compatibility across its plans at the same time as the iPhone 11 range launched in late September, allowing its customers to forego the regular physical SIM card. Optus and Vodafone already adopted the technology earlier in the year, allowing any compatible device to use its eSIM (embedded Subscriber Identity Module) in place of the regular plastic card that you'd otherwise need to insert. This also allows users who own devices compatible with both dual SIM and eSIM technologies to have two services active at once. This can be handy if you need to switch between numbers for work and personal use, for instance, or when travelling abroad and using an international SIM. It's not just the smartphone world that's been incorporating eSIM tech however – a wide range of LTE-capable tablets, 2-in-1s, laptops, and even smartwatches have launched with the new tech, and the list of compatible devices is growing exponentially. With that said, not every eSIM device is supported by Australian telcos, so it's worth checking with your carrier before doubling down.

Quality entertainment around the clock
Enjoy great entertainment throughout the entire home network.
Keep everything as simple as possible.



**GERMAN
DESIGN
AWARD
WINNER
2018**



AVM Fritz!Box 7590
Made in Germany



**Encrypted communication
via VPN**



**At home anywhere with
MyFRITZ!**



**Answering machine with
voice to mail**



**Ready for the future with
IPv6**



**FRITZ!Apps for iOS
and Android**



**Smart WiFi technology
Multi-User MIMO**



**Two wireless networks
for more flexibility**



**Band steering:
smart band selection
convenience**



**Mesh WiFi with FRITZ!
enjoy full mesh networking**




**Connected in no time
with WPS (Wi-Fi Protect-
ed Setup)**

Wi-Fi 6 certification now official

UP TO 3X FASTER THAN 802.11AC

Wi-Fi 6 is the new standard for everything that uses Wi-Fi – from phones and tablets to routers and desktops. As the new standard becomes commonplace, Wi-Fi connectivity will be faster and more efficient, especially in areas with high traffic. With the certification now official, companies are now free to have their Wi-Fi 6 compatible devices officially certified under the new standard.



Next-gen Vive Cosmos launched

BETTER SCREENS AND COMFORT.

The second generation of VR HMDs are arriving. Valve's Index was the first – which we define as a substantial increase in resolution, primarily – and HTC has just launched its Cosmos, featuring 2880x1700 combined res, which is an 88% increase over the original HTC Vive, as well as new LCD panels and a flip-up visor. The Vive Cosmos sells for \$1,299. We'll have a full review soon.



Chinese deepfake app lets you be in your favourite movies

BUT RAISES PLENTY OF PRIVACY CONCERNS AS WELL.

Chinese face-swap app Zao became an overnight sensation in its home country, but a clause in its user agreement raises privacy concerns. The app allows users to take a selfie and put themselves into a video – movies, shows, music videos or operas – but, according to the original agreement, the developers hold the rights to any imagery created using the app. Chinese journalists raised the issue and the app's user agreement has since been updated, although some people are still worried the app could be used for nefarious purposes.



Nvidia underperforms in laptops, report claims

FALLS WELL SHORT OF ITS DESKTOP EQUIVALENT.

A new report claims that Nvidia's GeForce RTX 2060 mobile GPU is a poor value proposition for laptop gamers, and that of all the RTX mobile offerings, it's considerably more off the mark compared to its desktop variant. This conclusion is drawn by **Notebookcheck.net**, which benchmarked the RTX 2080, 2070 and 2060 mobile graphics solutions in 3DMark, comparing them to their respective full-fat desktop graphics card versions.

Going by the 3DMark scores, the RTX 2080 mobile GPU was between 8% to 10% slower than its desktop version, with the difference being a more pronounced 14% to 18% with the RTX 2070. However, the real eye-opener was the RTX 2060 mobile which was 20% to 25% slower than its desktop counterpart.

The problem here being that when a consumer sees the identical names 'RTX 2060' for both desktop and mobile GPUs, they could consider them to be essentially the same product, with roughly the same performance levels.

Huawei first with integrated mobile 5G SoC

LEAPFROGS THE COMPETITION.

At Berlin's IFA show last month, Huawei announced its new Kirin 990 5G processor, along with a second non-5G Kirin 990. The 5G version is the first mobile processor that includes all key functions like CPU, GPU, NPU, plus memory and USB control – and now a 5G modem – on a single chip.

The Kirin 990 5G is a true unified design, supporting Sub-6 GHz 5G networks on both SA and NSA architectures. In order to keep the die size in check, Huawei is using TSMC's latest 7+ manufacturing process with EUV, which helps enable a smaller die size for the sorts of devices this chip will be going into.

To date, neither Qualcomm, nor Samsung, (nor Apple), have a unified flagship chip design that is near commercialisation. We do expect them to release the hardware as they generationally update, but as of today, Huawei is the first to announce it. These new chips are expected to be spread liberally across both Huawei and Honor flagships for the rest of 2019 and into 2020. ■

The AORUS logo is located in the top left corner. It features a stylized eagle head icon to the left of the word "AORUS" in a bold, italicized, sans-serif font.

WORLD 1ST GEN4 TOTAL SOLUTION PROVIDER

AM4 X570 Motherboards | Radeon™ RX 5700 Series Graphics Cards
AORUS NVMe Gen4 SSD | AORUS Gaming Monitor



<https://au.aorus.com/>

AMD
RYZEN | RADEON

GIGABYTE™

techbrief

» A DEEPER LOOK WHAT'S TRENDING IN TECH.

Chip yields, binning, and the silicon lottery

Jarred Walton explains why the current move to a smaller CPU process is important.

The primary building block of modern computers is the silicon microchip, made from 300mm silicon wafers, cut from a large silicon crystal cylinder. The monolithic crystals are grown from a seed crystal dipped into a molten vat of nearly pure silicon.

That "nearly" is important. Small impurities in the wafers can lead to errors or non-functional parts. It's not just impurities, though; the crystalline structure of the wafer also varies slightly. Chips from the centre of a wafer tend to be "better" than those near the edge, so they may require less voltage to hit the same clock speed, or may run at higher clock speeds at the same voltage.

Each silicon wafer contains potentially hundreds of chips, but not every chip from a wafer behaves the same. After the initial manufacturing process, each wafer gets cut into chips. These are tested to determine how good each is in a process called binning. Some may work perfectly, others may be effectively useless, while many fall somewhere in between. The percentage of functioning chips is the yield. Most companies don't report their yield, though rumours and estimates suggest that Intel's first-gen 10nm process had yields in single digits: Less than 10 percent of the potential 10nm processors were viable – and that's after Intel disabled the Gen10 graphics, and with a small dual-core die. In contrast, it's estimated that by harvesting partially functioning dies, yields for AMD's new Zen 2 parts may be 85 percent or higher.

This is important as each wafer has a relatively fixed cost – perhaps \$10,000–\$15,000. This is also why smaller chips are preferable to larger ones, as more can fit within the space of a wafer. A small 74mm² chip like Zen 2 has a potential of nearly 800 chips per wafer, while Intel's 10-core Skylake-X chips are around 322mm² and only fit about 170 per wafer. With a defect density of 0.1 per cm², the 10-core Skylake-X might lose up to a third of the potential chips (without harvesting), while the Zen 2 CCD would

only have about 8 percent defective chips.


Binning addresses multiple aspects.

Firstly, there's the harvesting of partially functioning dies – so your RTX 2060 that uses the same TU106 GPU as the 2070 only has 30 out of a potential 36 SMs enabled, as well as two memory controllers disabled. Binning also determines the ideal voltage and frequency, and sometimes a chip that is otherwise fully functional might get downgraded because it can't run at the desired speed. AMD's RX 5700 and RX 5700 XT use the same Navi 10 GPU, but besides having 256 fewer cores, the 5700 is

clocked about 150–200MHz lower.

The best chips are sold as the fastest, most expensive parts, while the functional but not quite as good chips are sold as lower tier parts. However, the gap between the "best" and "worst" functional chips from a wafer may be quite small. But for newer manufacturing nodes like TSMC's 7nm or Intel's upcoming 10nm, the gap is more likely to be larger. So, you won't hit the same clock speeds on second or third tier products, no matter how hard you try to overclock. Unless you're lucky, and win the silicon lottery with a great chip sold as a second-tier part. ■

"After the initial manufacturing process, each wafer gets cut into chips. These are tested to determine how good each is in a process called binning."



Intel's 10nm process has changed significantly to improve yields for the upcoming Ice Lake CPUs.



The home of technology

techradar.com

gadgets

» NICE LITTLE EXTRA THINGS TO HAVE

WD BLACK P10

Stick it in your Xbox.

\$138 (2TB), \$258 (5TB) | WDBLACK.COM

As games demand more and more bits, consuming precious console storage (*Cyberpunk 2077* will be at least 80GB!), eventually you're going to be faced with deleting old favourites games or files to make room for the new. Game drives have grown in popularity as a simple solution. Seagate has pushed hard into this market, and WD is offering up its new Black P10 for the Xbox One. It looks slick with its industrial styling, reminding us of a tiny shipping container. Performance is a bit limited and will see around 130MB/s at its peak. It can also be used on a PS4 or PC. A decent three year warranty is provided.



TEAMGROUP PD400

Everywhere with the data.

\$69 | TEAMGROUPINC.COM

Specialist solid state memory outfit TeamGroup is offering a reasonably priced SSD alternative to external portable hard drives. We've been testing the 480GB version, and its \$69 price would buy you around 1TB of external hard drive. Where you win is on speed. It's a regular SATA drive running through USB Gen 1 and will get you to around 430MB/s in transfers. It's also a tough little bugger with a sturdy case and proofing against shock, pressure and water. Some parts of TeamGroup's marketing say 'water repellent' others that it's 'water-proof'. We didn't submerge it in testing but there's no doubting its rugged materials and build.

D-LINK MINI WI-FI SMART PLUG

Smarten up your dumb old devices.

\$39 | DLINK.COM.AU

Plug anything into one of these little Wi-Fi power plugs and it's now welcome to join your burgeoning smart home family of connected things. The hub is the D-Link app, where on/off commands are issued, as well as schedules and routines are set. A new connected device is set up for the first time via the app and it's as easy as can be. Plug it in, name it and off you go. It integrates with Google Home and Alexa, so anything connected can be added to your preferences and settings, or added to any IFTTT configs. Connectivity is the older slower 802.11n, but that's all it needs. Good job, too, D-Link, for the slim design that won't get in the way of other plugs next to it on a board.





LOGITECH K580 SLIM MULTI-DEVICE WIRELESS KEYBOARD

Perfect for a particular person.

\$79 | LOGITECH.COM

The trick here is that this wireless Bluetooth keyboard has a switch that lets you change between different bound devices. The idea is that you pair your phone via Bluetooth, and the PC via the included Wi-Fi dongle. The phone rests in the cutout area (portrait or landscape) and when you're doing actual work on the PC and your friend texts you, just hit the button and respond to the phone text via the keyboard. Convenience! Reasonable value convenience, too, at just \$70, but as far as keyboards go it's poor, with keys and action reminiscent of old laptops before chiclet and scissor keys got really good, so this is not really for eight hours a day use.

TIVOLI AUDIO MUSIC SYSTEM

Retro. It never gets old.

\$1,199 | TIVOLIAUDIO.COM.AU

There's a lot to like about this fully featured smart speaker, but the eye catching retro styling is top of the list. Made with wood and fabric (and metal, and plastic...), it's reminiscent of the sort of gear that was popular in the 1950s, but with just about every new service supported. Via Wi-Fi or Bluetooth it'll stream whatever you want streamed, though we're most impressed that it includes a DAB/DAB+ digital radio – as well as good old AM/FM. Alexa is supported by not Google Home. There are four drivers inside, which goes part way towards explaining its relatively high price, but it really does pack almost everything in and looks lovely, too.




FITBIT VERSA 2

There are even more perks in Versa's V2.

\$329 | WWW.FITBIT.COM

Fitbit's most approachably priced smartwatch has been updated to pack in even more features without inflating the price. Headlining the Versa 2's list of updates is a new onboard microphone and Alexa integration that makes it easier to quickly get directions, search the web, or control the device hands-free. It'll also come with the latest developments in Fitbit's sleep software and introduces a new smart wake feature that sounds your alarm in the optimum part of your sleep cycle. A bigger AMOLED screen, faster processor and the option for an always-on display, have also been added to the long list of legacy perks continued from the original. ■



Are high end desktop platforms still relevant?

How many cores do you really need?

In years past, high end desktop (HEDT) platforms really stood out. More cores, higher clocks, more PCIe lanes, more memory channels and better expandability. Many of these traits are still apparent. But, now that AMD is set to offer a 16-core AM4 CPU, and Intel will soon offer a 10 core model on its upcoming mainstream platform, the question is, does the high end desktop have the appeal it once did?

HEDT will soon come back into the spotlight with the upcoming launches of some tasty gear. On paper, the HEDT platforms are attractive, no doubt. AMD's upcoming Threadripper 3 looks particularly exciting, especially if we get a 64-core variant along with its PCIe 4.0 support with a gazillion lanes and four (or more) memory channels. No doubt such a model will cost a trailer load of cash. Though not as feature packed or (r)evolutionary as Threadripper, Intel's upcoming Cascade

Lake-X platform should also generate some interest, though it remains to be seen just how competitive it will be in performance, features and importantly, price vs. Threadripper 3.

The HEDT platforms face challenges from their own manufacturers. AMD has the Ryzen 9 3950X 16 core with its high 4.7GHz boost clock. In all but highly scalable parallel workloads, it should beat out the high core count HEDT models while using less power. The same can be said for Intel's i9-9900K and 9900KS. Their eight cores and aggressive turbo modes mean their performance cannot be matched by the LGA 2066 parts under anything other than heavily multithreaded loads. They're also a lot cheaper, with the CPUs themselves and their supporting motherboards and memory all adding up to a significant saving.

Ask yourself the question, do you really need 32- or 64-core parts? If

you're a professional working day-in-day-out with large data sets and time sensitive highly parallel workloads, then of course you'll love them. Time is money. For non-professional users though, even those with software that can make use of 16 cores and 32 threads, is a 3950X enough? Will a workload that takes 40 minutes on AM4 vs 20 or 30 on a HEDT system really make that much difference to you? If it's your livelihood, then the answer will be yes, but for a hobbyist or prosumer? Probably not. Gamers will see no benefit from moving to HEDT unless they're limited by the number of PCIe lanes the mainstream platforms offer, particularly the Intel ones which struggle to run without compromise when running several NVMe drives.

HEDT was always niche, but now more so than ever. If you absolutely must have more than 16 cores, plan to use multiple expansion cards or expansive NVMe storage subsystems, then HEDT is for you. For the other 98% of users, the mainstream platforms are more than sufficient. HEDT is more like workstation lite. Motherboard manufacturers will still try to target gamers as much as ever with RGB bling and snarling warrior beasts, but it's clear that HEDT in 2019 is for professionals and content creators. For almost all other users, the HEDT is overkill with a capital O. ■

"Now that AMD is set to offer a 16 core AM4 CPU, and Intel will soon offer a 10 core model on its upcoming mainstream platform, the question is, does the high end desktop have the appeal it once did?"

Online anger may be silly, but the Epic Games Store is undeniably bad

If Epic wants to pour millions into exclusivity, it should also invest more in its user experience, writes Shaun Prescott.

If the internet has proven anything, it's that a sizable portion of games enthusiasts simply must have their voices heard. So it's no surprise that the new Epic Games Store is the target of increasingly visible online ire: the store doesn't have forums or any other social functionality. All the fury directed at the games store, cannot be contained within the games store. Twitter, reddit and Steam are awash in it, though.

Hell, the Epic Games Store didn't even have a search function when it launched. While the anti-Epic grievance brigade may seem outsized and, frankly, childish to the outsider at this early stage it would be impossible to make the argument that Epic's store is "good". It unquestionably sucks. It lacks functionality. It's yet another client to install. And Epic nowadays isn't even a particularly likeable company, considering it's crunch approach to *Fortnite*'s games-as-a-service development.

Not to mention that game enthusiasts with a collector's mindset have their collection fragmented yet again. This fragmentation is inevitable and growing: most major publishers have their own PC desktop clients, with Rockstar Games among the most recent to launch one. Steam's desktop monopoly may be slowly coming undone, but it still has a unique position and its client is undeniably feature-packed.


Epic Games' weapon in this war has been the acquisition of (often timed) exclusives, and the most recent exclusive title to release at the time of

writing is *Borderlands 3*. Gearbox's loot shooter is one of the most anticipated games of the year, but it's also the kind of game that requires ongoing patchwork, and Gearbox is definitely not a studio known for releasing bug-free day one builds. Traditionally, Steam's forums (not to mention forums hosted by other publisher clients like Origin, UPlay and Battle.net) have been a great place for PC users to flag bugs or other issues. But for lack of similar functionality on the Epic Games Store, PC gamers have resorted to, um, Steam. In the days following *Borderlands 3*'s September release, the Steam forums for *Borderlands 2* – which remains on Steam – were flooded with bug reports, complaints, and grievances related to the newer game.

If Gearbox head Randy Pitchford is to be believed, *Borderlands 3* had a very successful launch on PC. He claims its

number of peak concurrent players were nearly double those at *Borderlands 2*'s launch. Whichever way you spin it, that's an enormous amount of players all converging upon one very threadbare client: don't forget that *Fortnite* is also accessed via the Epic Games Store, and as of March 2019 that game had 250 million players (keep in mind though, that it's also available for smartphones and consoles – not everyone accesses it via Epic's PC client).

While it can manifest in some utterly stupid and occasionally frightening ways, you've got to understand the growing ire for this storefront. Epic removed its client development roadmap in August, meaning, as far as anyone knows, it's going to suck for a very long time to come. When you're working at this scale and have billions of dollars of *Fortnite* lucre to draw from, there really isn't any excuse. ■



END USER

Share your stories!

If you have an interesting story about technology users, their experiences and the issues that affect us all (whether funny or serious), email us at apcmag@futurenet.com. All correspondence becomes the property of APC and is subject to editing. Letters must include writer's full name, street address, suburb, state and phone number to be considered for print publication. Address and phone details will not be published.

» JOEL BURGESS STEPS INTO THE FRAY.

Is gaming the new black?

OLED emerges on the latest laptops, and Joel Burgess wonders if this dark horse display is capable of meeting the demands of creatives and gamers?

OLED is one of the most widely used display technologies in the premium TV space and AMOLED is the foundation for basically every smartphone screen, but other than a short-lived appearance on some expensive ultrabooks in 2017, the technology hasn't really been seen in PC displays... until recently.

Samsung's SDCA029 15.6-inch OLED laptop panel was shown off at the start of 2019 for CES, boasting a 4K resolution, 0.0005 to 600 nit brightness range and a 120,000 to 1 dynamic contrast ratio. This panel was also privy to the surprisingly low response times of OLED displays and benefits from wide viewing angles comparable to that of LCD IPS panels. What seems most critical this time around however is the cost, with these OLED panels not being any more expensive (and even at times cheaper) than identical devices with a UHD touchscreen display.

After testing the Gigabyte Aero 15 OLED, the Dell XPS 15 OLED and Asus' ZenBook Pro Duo OLED laptops (which all run the Samsung SDCA029 OLED) we've had a good chance to see what works in the first generation of OLED gaming laptops.

First off, this panel looks stunning as a laptop display. When watching movies and browsing the web the screen pops with colour, and even just the background and Windows theme have a vibrancy that is hard to unsee when returning to backlit LED displays.

There have been some issues raised online about the colour consistency with multiple sources pegging colour-accuracy variations 12 times higher than what's required for creative professionals. These concerns seem valid for the devices they were originally raised about (Alienware's m15 and Razer's Blade 15), but there are also a number of vendors that have used independent certifiers like the Video Electronics Standards Association (VESA) to independently test their screens HDR colour reproduction capabilities. Both the Zenbook Pro



"Samsung ditched pursuing OLED televisions in 2014 because of the problem of burn-in."

Duo and the Aero 15 OLED have DCI-P3 colour gamuts validated by Pantone and the Aero 15 has got and additional X-rite Pantone certification for a DeltaE colour accuracy of less than one, at 100% brightness, which makes it perfect for anyone needing colour precision.

Samsung ditched pursuing OLED televisions in 2014 because of the problem of burn-in – where blue sub-pixels degrade faster than the red and green sub pixels, causing distinct artifacts on any highly used areas of the screen. There have been reports of burn-in on in-store display models of LG's OLED TVs from 2017 after only 18 months, and people were even reporting it occurring on Google's Pixel 2 XL screens in the week following its release. To combat this Samsung's OLED laptops have blue pixels that are twice the size of their green and red

counterparts, which will allow them to be less bright and, hopefully, last longer.

The direct pixel activation in OLED means that the display doesn't have to work out the best backlight setting for the entire picture before switching. This gives the Samsung OLED panel a sub 3ms response rate ... which is about four times faster than Sharp's 240Hz ultra-fast gaming laptop panel. Yet Samsung's OLED panel is only available in 60Hz, which will be a drawback for anyone hoping for a smooth competitive-level gaming experience. That said, even if you get an OLED laptop with an RTX 2080 (Max-Q) GPU, you'd struggle to get more than 60fps on UHD Ultra settings, so there's no real need for high refresh 4K OLED on a laptop just yet. ■

numbercrunch

» WE LOOK AT THE NUMBERS DRIVING THE BIG TECH NEWS.



18945

The build number for a new edition of Windows 10 currently being tested by Insiders. This build shows the future of Cortana as a chat app, rather than integrated in the Search box.

ONE THIRD

The number of fitness tracker owners who stop using their device after six months, according to consulting firm Endeavour Partners' survey.

FORTNITE

US\$3,000,000


The prize pot taken away by the winner of the Fortnite World Cup held at the end of July. Total prizes in the tournament reached US\$30 million.

04

The number of GPUs Intel appears to be readying for market. Little is known about the new Gen 12 graphics chips, except that, with up to 512 execution units, the chips, known as Xe, seem to be aimed at the likes of Nvidia and AMD add-in boards rather than the usual integrated graphics from Intel.

\$1.45 BILLION


The cost to Apple of buying Intel's modem-making business. Think sleek smartphone 5G chips, not buzzy 56K PCI models.



ONE BILLION

The number of connected wearable devices – fitness bands and the like – worldwide by 2022, according to figures from Statista.

how it's done



Huawei Mate 20 X 5G

Peeking inside its great walls.

While the US-China trade dispute drags on and we ponder Huawei's future, the company charges into the 5G arena with its new Mate 20 X 5G. We're curious to see what this plus-sized phablet conceals, and how independent the hardware might be – even as Huawei faces the possibility of being banned from doing business with US companies.

MAJOR TECH SPECS

- 7.2-inch 1080×2244 OLED multitouch display
- Octa-core Huawei Kirin 980 chipset, with 8GB RAM
- 256GB onboard storage
- Balong 5000 multimode 5G modem
- 4,200mAh battery with 40W SuperCharge 2.0 support
- Triple rear cameras: 40MP f/1.8, 20MP f/2.2, and 8MP f/2.4 lens with 5x optical zoom

KEY FINDINGS


- On the bottom we recognize the usual suspects: a USB-C port, two mic holes, and a speaker grille. Along the top, we find another mic hole, the infrared blaster, and a sliver of a grille for the earpiece speaker. Compared to the already-large Mate 20 Pro, the X 5G looks colossal. On the back, we note

the fingerprint sensor beneath the camera array. Although this Mate is only rated IP53, the SIM card tray is equipped with a rubber gasket – something we usually see on “waterproof” smartphones. To our surprise, we get the glued-down back cover off lickety-split – no heat required.

- A generous fingerprint flex cable keeps the back cover tethered for now, but it's so long that we don't mind. A flock of screws hold the NFC coil, antenna, and graphite heat conductor pad in place. With those out of the way, we can disconnect the fingerprint sensor and get our first good look inside. The 24MP front-facing peeper comes out with a simple pry. We'll never grow tired of repair-friendly press-fit connectors. The motherboard also comes out easily, allowing us to unplug the three-eyed rear camera block.
- It looks like the 5G modem comes with its own block of LPDDR4X memory. To access the battery, we remove the board interconnect cable. We also pry up the glued-in loudspeaker, and the daughterboard with the USB-C port.
- We follow the built-in battery removal instructions step by step.

This is the same battery used in the Mate 20 Pro, weighing in at 16.04Wh. That's a monster compared to the iPhone XS Max's dual-cell 12.08Wh powerplant.

- The display adhesive holds on tighter than the back cover adhesive, but gives way after some heat and halberd slicing. There's no fancy under-display fingerprint sensor, just a blank OLED screen and the aluminum frame. As in the standard Mate 20 X, there's a big vapor chamber hiding on the aluminum frame behind the graphite foil.
- Repairability score: 4 out of 10 (10 is easiest to repair). Many components are modular and can be replaced independently. The battery is accessible after removing the rear cover and midframe, and has pull tabs. Standard Phillips screws are used along with an average amount of adhesive. The midframe covers the battery and fingerprint sensor connectors, and is obstructed by camera sensors and flash. Glued-down front and back glass means more risk of breakage, and makes all repairs difficult to start. Screen repairs will require near-complete disassembly. ■



Find the best prices
on the latest tech.
www.getprice.com.au

GETPRICE  TM



CRITICAL SPECS:
Windows 10 Pro 64 bit;
15.6-inch 4K OLED
display @ 3840 x 2160
pixel resolution;
ScreenPad Plus
resolution 3840 x 1100;
Intel Core i7-9750H or
i9-9950HK CPU; 32GB
of RAM; Nvidia GeForce
RTX 2060 GPU; 1TB PCIe
SSD; 73Wh battery
(2.5h 1080p movie
playback); 35.9 x 24.6
x 2.4cm; 2.5kg.

LAPTOP

\$4,999 | ASUS.COM/AU

Asus ZenBook Duo

Is this dual screened workstation the one for creative professionals?

What's better than a 4K screen on a professional ultrabook? One and a half 4K screens on a professional ultrabook. As you might have guessed, Asus's latest brainchild, the Zenbook Pro Duo, has two screens: a 15.6-inch Full HD or 4K OLED panel in the usual spot, and a second, half-height 14-inch touchscreen that sits directly underneath it, just above the keyboard, with a resolution of 3840 x 1100.

While it's logical to think that Asus decided to take Apple's idea of a touch bar and expand it into a full touchscreen, the catalyst for the ZenBook Pro Duo's additional screen, we've been told, was actually just one of opportunity. Back on the first Zephyrus, Asus

relocated the trackpad to the bottom right of the device (where it doubled as a digital number pad), which left a big gap above the keyboard. After a couple of years Asus decided it was time to put this dead space to use as a second LED touchscreen to help boost productivity on the go.

We were a little apprehensive about just how essential a second touchscreen was, especially when you factor in the additional battery drain. The half height IPS touch panel runs at the same FullHD or 4K horizontal resolution as the main display to keep everything in line, but it isn't quite as bright and has a matte finish that makes it look much duller than the OLED

above it. The ScreenPad Plus, as Asus calls it, is recognised by Windows 10 as an additional display, allowing you to extend the screens and arrange them vertically, rather than just being an expansion of the one desktop space.

The ScreenPad Plus is a multi-touch panel that lets you interact with it using your fingers or with the accompanying stylus. To expand its potential, Asus custom-built an app draw that includes a touch number-pad, recently used command hotkey bar and a handwriting word recognition app. So all up we're pretty impressed with this second screen.

The ScreenPad Plus features all the Windows 'Snap' resize functionalities

you get on additional screens, but it also has further customisations allowing you to quickly partition the screen into three small screens, which is good for ultra wide 32 by 9 form factors. There are also integrations with prolific music, photography, programming, video editing, game streaming and animation software to make sure the toolbars and feature windows work well in the ScreenPad display. Even if you just have multiple web browsers open, the ScreenPad Plus makes split screen multitasking much easier on a laptop.

We've talked a lot about the second display, but what's perhaps even more impressive is the main 15.6-inch 4K OLED that sits above it. It's the same Samsung OLED panel that was announced at CES earlier in the year, but Asus has done a lot of in-house calibration to get one of the widest colour gamuts we've seen on an OLED laptop.

"We were a little apprehensive about just how essential a second touchscreen was, especially when you factor in the additional battery drain."

"The ZenBook Pro Duo is pitched more as a creative workstation for professionals than it is a gaming laptop, but with a Nvidia RTX 2060 GPU, it can pretty confidently do both."




LABS BENCHMARK RESULTS


PCMARK 8 - HOME (SCORE)




CINEBENCH R15 - CPU (MULTI-THREADED)




HWBOT X265 - 1080P (AVG FPS)




3DMARK FIRE STRIKE (SCORE)




GHOST RECON WILDLANDS - ULTRA 1080P (FPS AV)




CRYSTALDISKMARK READ (MB/S)



CRYSTALDISKMARK WRITE (MB/S)



PCMARK 8 - HOME BATTERY ACC. (HH:MM)



Asus is boasting 100% DCI-P3 and has certifications from VESA for independent True Black 500 HDR and a colour calibrated profile from Pantone. This additional colour potential blended with the true black of the OLED panel and the slim 5mm bezels makes for an impressive visual display that is great for creative work (especially cinematography, since DCI-P3 is a cinematic colour gamut).

The Zenbook Pro Duo UX581 can be configured with either an Intel Core i7-9750H and 32GB of RAM for \$4,799, or with a Core i9-9980HK and 32GB of RAM for \$4,999, but only the latter was available for pre-order at the time of writing, with shipping expected to begin on the 25th of October. We tested the Core i7 variant, which uses a hexa-core chip with a base frequency of 2.6GHz that'll boost to 4.5GHz. The chip performed as well as

you'd expect from a top shelf consumer laptop processor, landing within 10% either way of the average scores we have for this CPU in general work and raw processing tasks. This chip regularly spiked to 100 degrees and while this didn't impact the performance, we do wonder if the Core i9 will show any improvement since it'll use the same thermal management system.

The ZenBook Pro Duo is pitched more as a creative workstation for professionals than it is a gaming laptop, but with an Nvidia RTX 2060 GPU, it can pretty confidently do both. In raw processing scores it was only between 2% and 5% ahead of an Aorus 15 running an RTX 1660Ti GPU. This is a bit lower than we'd expect, but it also happened to net between 9% and 23% off the average frame rates of an Aero 15 OLED running an RTX 2070 on our real-world gaming benchmarks. While the

2060 is powerful enough to ensure 30fps plus framerates at 1080p ultra settings, if you do want to game in 4K you'll have to play older games or use lower graphical settings – which makes it hard to use this screen to its full potential for gaming. Fortunately the HDR colour profile is still detectable at Full HD and the OLED panel looks vibrant at any resolution.

A high resolution screen requires a lot of energy to maintain and while the ScreenPad Plus is not a particularly bright display, the ZenBook Pro Duo is still running 50% more screen real estate than its competitors. From this standpoint it makes sense that the 73Wh battery on the ZenBook Pro Duo only lasts two hours and 13 minutes in the PCMark 8 Home (accelerated) battery life benchmark, 47.5% less than the average of the other OLED devices we've tested. It lasts slightly

longer during 1080p movie playback netting two hours and 51 minutes, just 25% behind the Dell XPS 15 OLED, but this is still on the low battery life end for professional laptops.

You'll also have to sacrifice some portability on the ZenBook Pro Duo since the laptop is between 7 and 4mm thicker (at 2.4cm) and at least 500g heavier (at 2.5kg) than the Dell XPS 15 or Gigabyte Aero 15 OLED devices. This isn't unwieldy though in the scheme of things. There's also a bit of getting used to the inline trackpad placement and front stage keyboard setting, but we didn't find it particularly disruptive.

■ Joel Burgess

Verdict

If you're happy to pay a premium, the ZenBook Pro Duo's second screen is a compelling productivity perk.



"If you're happy with 1080p gaming the Nvidia GTX 1050 GPU is serviceable, netting around 30fps for current games on Ultra 1080p settings."



LAPTOP

FROM \$2,199 | WWW.DELL.COM

Dell XPS 15 OLED

Dell goes after the detail-oriented with its latest XPS 15 OLED professional ultrabook.

The new Dell XPS 15 can be configured with an overclockable Core i9-9980HK CPU, 32GB of RAM and a 1TB NVMe SSD in RAID 0 configuration for \$5,199. This is a powerful CPU that we've really only seen on overpowered gaming rigs and while it's great to see this chip on Dell's ultrabook, it requires a generous cooling infrastructure to work well, and we're not fully convinced that Dell has nailed the thermals just yet.

The real headline feature, however, is a new 15.6-inch 4K OLED screen model that is 20% dimmer than the more expensive UHD touch display, but offers amazing 100,000 to 1 contrast ratio and rich dark colours. It's claimed to include 100% Adobe RGB colour profile and 100% DCI-P3, depending on which part of Dell's website you look at.

Full Adobe RGB would make it the broadest colour gamut of the laptop OLED screens we've seen to date, however you'll have to take Dell's word for it since it doesn't offer any third-party calibration or HDR certifications. Notebookcheck measured the screen's capability at 81% AdobeRGB 1998, which is shy of Dell's claims but still much better than the 68% scored by the Gigabyte Aero 15 OLED.

While this device does offer Dolby Vision HDR movie playback, the screen output seems to be locked to the Intel UHD Graphics 630, so you won't be able to game in HDR, even though the screen and GPU are both capable of it. It's not uncommon for laptop vendors to configure GPUs in this way in order to minimise battery expenditure when they're not in use, but it does isolate a

number of GPU functionalities (including HDR gaming) from the screen.

We tested the Dell XPS 15 with an Intel Core i7-9750H CPU and 32GB RAM allocation that is listed on Dell's website at \$4,299, but we've seen on sale for \$3,654. While raw CPU and media encoding scores of the XPS 15 were between 10 and 17% better than our average for this CPU, the general work tasks were between 3 and 20% behind other comparable devices. This is potentially due to some throttling over longer CPU-heavy tasks since we were regularly able to get the CPU to spike to 1000C. The i9 model, presumably, uses a similar thermal infrastructure to the XPS 15 i7, which makes us a little worried about the potential for the i9 to offer any performance benefits at all.

If you're happy with 1080p gaming the Nvidia

GTX 1050 GPU is serviceable, netting around 30fps for current games on Ultra 1080p settings. If you're playing slightly older games or you're happy to compromise on graphical settings then you can even get some titles running at 30fps in 4K.

Despite a large 97Wh battery, the device only lasts four hours and seven minutes on the PCMark 8 battery benchmark and less for movie playback, which means it'll be a bit shy of the full day mark if you use it sparingly.

■ Joel Burgess



Verdict

A great option for professionals that need printable consistent colours and a little extra graphical processing.





LAPTOP

\$2,199 | ACER.COM

Acer Predator Helios 300

This entry-level gaming laptop ticks all the boxes, but we're not sure it has the mass to make it really gravitating.

Acer seems to have done some work on its thermal interface for the Predator Helios 300, introducing some cool looking blue radiator grills, we assume inspired by Dell's G5 15 Special Edition, and a nuanced rear vent shape that is borrowing some elements we loved on the past few years' worth of Aorus Laptops. Unfortunately these appealing details are hampered by a predictable square top edge and a lightweight plastic undercarriage, which drags the overall aesthetic back into the mundane.

The Helios 300 15-inch and 17-inch models are squarely aimed at the entry level gaming market so none of their specs are going to seem out of the ordinary. Both units sport an Intel Core i7-9750H CPU, 16GB of RAM and a 512GB PCIe SSD. Where they differ is in the screen and the accompanying GPU, with the former netting the

more recent GTX 1660Ti for an RRP of \$2,199 while the latter has the more powerful RTX 2060 and goes for \$2,499.

In addition to size the screens also have different refresh rates, with the bigger screen getting the fast 144Hz display and the 15-inch model getting just a 60Hz refresh rate. We benchmarked the smaller model with the 15.6 inch IPS panel, but both should actually perform closely, save for the highlighted differences.

The Intel Core i7-9750H CPU is a pretty solid choice for powerful laptops and the Helios 300 device we tested generally performed better than our average for this chip in general work and raw CPU stress benchmarks, but it was still well within the 10% variance buffer you tend to see between CPUs on these benchmarks. This CPU offers more than enough performance for the average

workstation user with the 2.6GHz hexa-core processor chewing through demanding tasks like 1080p media encoding at 33 fps.

The Predator range is pitched directly at the gaming market and while Nvidia's newish GTX 1660Ti is still an entry-level gaming card, it offers exceptional value for money. The performance across all our real world gaming benchmarks was near identical to the Aorus 15 we tested recently, with each average within a couple of frames of each other. For games like *Metro: Exodus*, *Ghost Recon: Wildlands* and *The Division 2*, you get framerates at 39, 45 and 53 respectively, so it's a solid 1080p gaming device. Since this screen tops out at a maximum of 60 Hz, it's pretty well suited to the GPU power offered here and it makes sense that the 17-inch 144Hz display of the Helios 300 is paired with the more

powerful RTX 2060 GPU.

The two hours and 31 minutes that the Helios 300 got in PCMark 8's Home battery accelerated benchmark is about half an hour more than what we pegged on the Aorus 15, but it's still well within the timeframe for a short-lived gaming laptop battery.

The keyboard is a little shallow and rubbery to be a delight to type on, but both it and the trackpad are perfectly serviceable. There's also enough interface options to mean that you shouldn't have any problems connecting this unit to your existing workstation accessories.

■ Joel Burgess


Verdict

Everything you would want in an entry-level gaming laptop, but it's missing anything that gives a 'star quality'.




LABS BENCHMARK RESULTS


PCMARK 8 - HOME (SCORE)




HWBOT X265 1080P (AVG FPS)




CINEBENCH R15 - MULTI-THREADED CPU




3DMARK TIME SPY (SCORE)




GHOST RECON WILDLANDS - ULTRA 1080P (FPS AV)




TOTAL WAR: WARHAMMER 2 - ULTRA 1080P (FPS AV)



CRYSTALDISKMARK READ (MB/S)



CRYSTALDISKMARK WRITE (MB/S)





LAPTOP

\$1,899 | WWW.ASUS.COM

Asus Zenbook 14 UX433F

The 14-inch laptop that's almost all screen.

The world's smallest 14-inch laptop. That's Asus's sock-it-to-you sales pitch for the new Zenbook 14 UX433F. It's not the only thing this slick Intel-powered laptop has going for it, of course.

There's powerful quad-core processing, just for starters. Claimed all-day battery life and Harmon Kardon audio sweetens the deal, too. But it's that slim-bezeled 'NanoEdge' display and the extremely compact proportions it enables for a laptop in the 14-inch segment that most defines the Zenbook's proposition.

What's on offer here, in effect, is a laptop that feels like it's from the segment below in terms of size, but delivers a larger, more ergonomic 14-inch viewing experience.

That means this 14 inch is barely any bigger than many 13-inch ultrabooks, like Dell's XPS 13, while 14-inch Ultrabooks including the Lenovo ThinkPad X1 Carbon are

tangibly larger. Whether that's enough to stand out in a seriously crowded market of premium portables, is another matter.

The Asus Zenbook 14 UX433F is all about that screen. Or rather the lack of anything much around it. The side bezels measure just 2.9mm, while the bottom clocks in at 3.3mm and the top 6.1mm.

The net result is a pretty staggering 92% screen-to-body ratio. Most notable is that slim bottom bezel. It's what allows the Zenbook 14 to match the depth of 13-inch Ultrabooks with relatively hefty chin bezels, like the Dell XPS 13, for overall depth, while only being slightly wider in width.

What the Asus Zenbook is not, however, is a truly ultra-thin and ultralight laptop. The likes of the aforementioned Dell XPS 13 is noticeably slimmer if no lighter and the super slim Acer Swift 7 is a 14-inch alternative in a

whole different ballpark for thinness, clocking in at 0.89cm to the Zenbook's 1.59cm.

That said, there are upsides to a little thickness, including space to fit a pair of full-sized USB A ports. It's just a pity that only one of them is full 3.1 speed while the other only offers the ancient and slow 2.0 specification.

If the Asus Zenbook 14 UX433F is all about that screen, it's a pity it isn't slightly more impressive beyond those headline grabbing minimal bezels. However, it's not a dud of an LCD panel, to be sure.

The performance of the other core components is more impressive. The Intel Core i7-8565U quad-core CPU will be plenty for all but the most demanding mobile workflows. Ditto the Intel 512GB NVMe SSD, which delivers desktop class storage performance.

If there are exceptions to all that, they involve the graphics subsystem and the mere 8GB of RAM. The

former is no surprise for an Ultrabook, just be aware that like any current laptop that utilises Intel's integrated graphics, it'll be pretty useless for all but the least demanding games.

The impact 8GB system memory limitation is less clear cut. Much will depend on your typical usage.


For all its slim-bezel goodness, the Zenbook's screen is also merely OK. It's not the highlight it would have been with a punchier, higher resolution panel. All of which makes the Asus Zenbook 14 UX433F a solid choice worthy of your shortlist but not necessarily a no brainer. ■

Jeremy Laird

Verdict

A great looking machine with reasonable performance and a very nice price.





MONITOR

\$799 (27-INCH), \$999 (32-INCH) | SAMSUNG.COM/AU

Samsung Space Monitor

Samsung is working hard to make monitors look good from behind.

Samsung had a lot of success with its Frame televisions, which sacrificed top-of-the-line specifications for a TV that looked as nice switched off as it does when in-use, so it seems only natural that its a stylish, minimalist monitor for the design-conscious PC user should follow.

The key component in Samsung's new Space Monitor is an easy to install support arm that clamps to the rear of your desk and lets you quickly adjust or swap the screen's position to suit your needs. On the one hand, the screen can be pushed right to the back of the desk, with the monitor arm standing straight up, even if the desk is flush against a wall. While this might seem like something that all monitors should intrinsically do, you're often either fighting against an expansive monitor base, a large rear curve (even on flat screen

displays), or a surprisingly bulky mounting system – all of which cause the screen to protrude from the back edge of the desk. At the other extremity, the arm can be brought right down to desk height using a second neck hinge that tilts it up to keep it facing you.

While this mounting arm is versatile and simple to use, it doesn't include the option for any horizontal rotation or screen tilt, so if you want to orient the screen vertically, use it on a corner of the desk, or try to use multiple monitors then you're going to have issues. It's also not compatible with any desks thicker than 9cm, or anything with built-in draws that might impede the mounting clamp.

The new Space Monitor comes in two sizes, a 27-inch 144Hz display at WQHD 2560 x 1440 pixel resolution, or a 32-inch 60Hz at 4K 3840 x 2160 pixel resolution. With a typical brightness of just

250 nits, the Space Monitors are quite a bit duller than the latest IPS displays and you definitely notice it in brighter environments.

The display is advertised to have a 100% sRGB colour gamut, which will be fine for non-creative workflows, but is much lower than what you want for any design work. Samsung is suggesting that the VA panels here have 1780 viewing angles, but we definitely noticed some discolouration at angles greater than 450.

While there's no G-Sync or Freesync tech built into either model, the smaller 27-inch display can output at 100, 120 or 144Hz, allowing you to play games at fast frame rates, provided you have a sturdy gaming GPU.

The Space Monitor is pretty streamlined when it comes to inputs, with just a Mini DisplayPort 1.2 and a HDMI 2.0 input. Accompanying these inputs is a clever cable that merges the power

input and a HDMI cable and feeds them down the back of the monitor arm with inputs at the base.

It's not hard to find better quality displays for less, so if you're willing to get the screwdriver out then paying the \$100-odd dollars for a good desk-clamp VESA monitor arm seems like a no-brainer. If however, you do care as much about what the rear of the monitor looks like as you do the front, then it's hard to think of a more suitable display.

■ Joel Burgess

Verdict

An expensive professional monitor that includes a simple to use monitor arm for a considerable markup.





NAS

\$750 | [HTTPS://WWW.SYNOLOGY.COM](https://www.synology.com)

Synology DiskStation DS620slim

Is this mini NAS as good as Synology's excellent bigger siblings?

Need a NAS in your home or office, but physical space to stash it is at a premium? Maybe you need to lug a NAS with you on the road? Synology's new DS620slim NAS might be the answer to your prayers. This little guy has six 2.5-inch drive bays in a compact 121mm x 151mm x 175mm rectangle, so it can fit in the tightest of comms cupboards, studio apartments or Pelican cases.

Powering the DS620slim is an Intel Celeron J3355 CPU, which has QuickSync support, allowing this NAS to do hardware encoding and decoding of H.264 and HEVC (H.265) video. Nifty for any Plex users out there. There's 2GB of RAM from the factory, but there's an SO-DIMM slot free for you to install extra memory if you need it.

The closest match in Synology's line-up for the DS620slim is the DS918+. Both sell for roughly the

same price (approximately \$750), but the DS918+ has a superior quad-core CPU, 4GB of RAM and support for 3.5-inch and 2.5-inch drives. It's difficult to imagine a scenario where the DS918's roughly 4cm of extra girth is such a burden that it justifies a slower CPU and less RAM in the DS620slim. Really, the DS620slim's only practical advantage are its two extra 2.5-inch bays.

For those using SSD only storage in their NAS, the DS620slim's extra two bays compared to the DS918+ mean you can purchase additional 1TB or 2TB SSDs for more storage, instead of having to shell out for expensive 4TB SSDs. For those who prefer speed and power savings on a budget over raw storage, the use of six 2.5-inch HDDs can provide a speed boost and smaller power bill in a RAID configuration, that four 3.5-inch HDDs cannot.

Synology's other six-bay NAS, the DS1618+, isn't that more expensive than the DS620slim (\$1,150 vs \$750), yet is much, much faster with its Intel Atom C3538 CPU and 4GB of RAM, plus way more expandable with two RAM slots, four Ethernet ports and even a PCIe 3.0 x8 slot that allows for NVMe PCIe SSDs. The DS1618's storage configuration is more versatile than the DS620slim's too, as the 3.5-inch bays also support 2.5-inch drives, allowing for a mix of sizes that the DS620slim can't provide.

The DS1618+ is a beefy unit though, at almost twice the size and weight of the DS620slim. So, if you carry your NAS around in the field, need six drive bays and are happy to sacrifice performance for compactness, the DS620slim will fill the storage related hole in your niche heart.

While the DS620slim is a solid product and full featured thanks to the lovely DiskStation Manager software, it occupies an extremely niche spot in Synology's NAS line-up. The situations where its dual-core CPU and 2.5-inch only drive bays provide better value than the DS918+ or DS1618+ are far and few between. Unless you're carrying your NAS with you regularly and need the smallest form factor possible at the expense of performance, it's hard to recommend the DS620slim over Synology's other units.

■ Anthony Agius

Verdict

If you need a NAS that's compact but still takes six drives, the DS620slim fills that niche.



"Impressively for the price, the access point includes the ability to create a mesh network of up to eight VigorAP 903s."

ROUTER

\$280 | DRAYTEK.COM.AU

DrayTek VigorAP 903

A professional level wireless access point with mesh technology.



Aimed at businesses who want to expand their Wi-Fi coverage, the DrayTek VigorAP 903 is designed to be an easy to setup and manage solution. At the time of writing the VigorAP 903 is quite new so there aren't a lot of options for shopping around, but the price is still quite competitive at \$280.

The access point uses AC1300 spec Wi-Fi, with two dual band radios and removable antennas. It can handle eight different SSIDs and two virtual LANs, to create different user groups with specific network rules. The AP also has a five port Gigabit switch built in, and supports power over ethernet. While the access point can also be plugged in normally, PoE is an important feature, as it means it can be installed in with just one cable. The VigorAP 903 also has a USB 2.0 port, that can be used to

drive a printer, or even monitor the temperature with a USB thermometer.

While the underlying hardware is solid, where the VigorAP 903 stands out is the software features. Impressively for the price, the access point includes the ability to create a mesh network of up to eight VigorAP 903s. We tested two units to get hands on with the feature, setting one AP as the root device via the web interface. It's then just a matter of powering up another AP in the desired location, then telling to network to add it as a new member using the existing configuration. In a larger network, this means you can quickly deploy an access point to fix a dead spot, or re-arrange the network. If one of the access points goes offline, the network automatically reconfigures to maintain the best performance. The VigorAP can have wired

LAN devices plugged in to connect them to the network, or link to other APs with a cable to improve performance. The VigorAP 903 can also operate as a standard access point, or range extender. It has all the in depth features you'd expect, such as load balancing, airtime fairness, band steering, assisted Wi-Fi roaming and scheduling.

Setup via a web browser, though there is also a mobile app and a Windows configuration program. The Android app is a big letdown, with instructions in broken English, and despite testing with multiple devices it simply would not connect to the access points. Fortunately the normal browser interface is excellent, and makes it extremely easy to configure and manage the routers.

With 867Mbps available

on the 5GHz band, and 400 Mbps on 2.4GHz, the router is aimed at creating a solid network, not outright speed. During testing we found it very dependable, but with fairly average performance. At close range we pulled 301 Mbps, and at 10m (with walls in the way) we managed an acceptable but not stand out 131 Mbps.


For a business that wants a robust yet easily managed wireless network, it's hard to go past the VigorAP903. It's affordable, easily expanded, and offers decent performance. ■

Lindsay Handmer

Verdict

Solid pricing, easy setup and excellent featureset easily make up for the staid Wi-Fi performance.





MEMORY

\$155 (3000MHZ), \$169 (3200), \$219 (3600) | THERMALTAKE.COM

Thermaltake ToughRAM RGB 3200

"OK Memory, twinkle like a little star."

Here's the latest soldier in the battle to change memory from a commodity to a luxury hero item. This particular design is, somehow, a little less show-offy than some of its competitors, with a relatively restrained set of coloured blocks on top in a fairly elegant and simple arrangement. I think it looks pretty decent, and I usually have this stuff tuned off. It hooks into Thermaltake's lighting management software TT RGB Plus, as well as Razer Chroma, combining forces with your other light sources within your PC to shout colour at your feet. And, because this is the future none of us asked for, it responds to voice commands through Alexa or TT's own app.

Thermaltake has non-RGB ToughRAM to come, though we don't know the final specs yet. The company has said that the non-RGB models will start at 2,400MHz and top out at 3,000MHz. Liquid cooling

variants haven't been announced. For now there are 3,000MHz, 3200MHz and 3600MHz RGB kits available, all 2x 8GB.

While the frequencies available put these kits on the radar for those seeking decent performance, the timings are unfortunately the weak link if you want that little bit extra. The 3200MHz kit we have here, as well as the 3000MHz kit, have XMP timings of 16-18-18-38. The 3600MHz kit uses 18-22-22-42. These are mediocre at best and in line with budget RAM speeds. In day to day use they'll hold up just fine, but this isn't a kit that anyone wanting to aggressively overclock their memory would be looking for. ■ Ben Mansill

Verdict

Unusually attractive RGB, but sub-standard performance due to poor memory timings.



ROUTER

\$279 | WWW.TP-LINK.COM/AU/

TP-LINK Archer MR600 4G Router

For when wired connections aren't a reliable option.

The MR600 is a (very) basic, old-school, no-frills Wi-Fi router with four Gigabit Ethernet ports (one of which is WAN-compatible) and not much else... except for a 4G MicroSIM slot. We took a SIM out of our phone, stuck it into the slot at the back with an adapter, entered the default Wi-Fi details and suddenly the router was sharing 4G internet on a Wi-Fi network. Speed tests are borderline irrelevant as you're limited by local phone network but maximum throughput is stated as 300Mbps downstream and 50Mbps up. In our office, where Telstra speed is reasonable by the street, (but there's no signal whatsoever at the far end of the office) the router provided a great amplifier for everyone attached to its Wi-Fi. Up to 64 devices can connect.

Managing the device is simple thanks to TP-Link's Tether smartphone app which enables you to change network names and passwords and combine

2.4GHz and 5GHz channels into one "Smart" channel that ensures devices are connected in the fastest way possible. There are also parental controls but not much else.

Compared to the small, nascent, 5G, personal Wi-Fi hotspots that the telcos push, this is less portable due to its lack of battery and larger size. However, as the basis of a small or transient office, with multiple people connecting to it - that has no option but to access 4G - the more-robust hardware and superior antennae make it a great choice. Just note that the \$279 price will be dwarfed by ongoing data costs. ■ Nick Ross

Verdict

A basic but functional 4G-based router.





MOUSE

\$89 | MICROSOFT.COM

Microsoft Pro Intellimouse

Hello, old friend.

The nascent stages of a retro revival appear to be underway. Recently Logitech released an updated version of its positively ancient MX518 mouse, and now, under my hand, sits the Microsoft Pro Intellimouse.

Compared to the gaming mice of today the Intellimouse is positively enormous – and that's the charm. It's just so damn comfortable. That lovely arch across the top supports the palm like no other mouse and is a very relaxing experience. I'd forgotten how nice this beautifully sculpted design is. This thing was a game-changing smash hit when it was originally released and much of the praise heaped upon it was for its then-revolutionary optical sensor, as much the ergonomics.

Thankfully the original 125Hz sensor has been replaced with a cutting edge 1,000Hz 200-16,000 dpi

PikArt sensor. It can be adjusted in increments of 50 dpi, too, so you're virtually guaranteed to find a sweet spot that's perfect for you. The button switches are also new, but have a satisfyingly old-school heft to the movement and actuation. Clicking feels like a decisive action, accompanied by a relatively loud clicking sound. I do like this.

It's wired, thus eschewing its last chance to modernise by going wireless like many new mice, but that's ok too. This is for people who want a sturdy and reliably mouse that's comfortable, doesn't shine through your hand with an RGB spectrum and reminds you of the good old days. ■ Ben Mansill

Verdict

Solid, dependable, frill-less and with a good dose of retro tech nostalgia.



MOUSE

\$249.95 | WWW.LOGITECH.COM/EN-AU

Logitech G502 Lightspeed Wireless Gaming Mouse

Wireless movement with wired speed and precision.

For serious gamers, especially those who play competitively, you can't go past a wired mouse for speedy, lag-free, precision tracking. Of course, the freedom of movement that a wireless mouse can provide is undeniable. Enter Logitech's G502 Lightspeed Wireless Gaming Mouse, a pointer that employs a lightning-fast 16,000 CPI HERO sensor for instantaneous travel between distances. With an IPS (inches per second) rating of 400, the G502 can move as fast as a wired mouse without stuttering or loss of tracking. Sporting a total of 11 programmable buttons, the G502 can be tuned to perfection for any gaming situation in Logitech's G HUB app.

That said, button placement will occasionally frustrate – the player's thumb is surrounded by incredibly-sensitive buttons that are very easy to press by accident. Still, the

buttons all feel nice and clicky, and we appreciate the scroll wheel's ability to toggle between free-spinning and chunky stepping.

It's also modular, with an included set of 2g and 4g weights allowing users to make the G502's chassis heavier by 16g. Players can also take advantage of the G502's fully-customisable Lightsync RGB functionality, which will also sync up with Logitech's PowerPlay wireless charging mats. At \$249.95, the G502 Lightspeed is very expensive, but when it comes to flexibility, speed and accuracy, it's easily Logitech's most advanced gaming mouse to date. ■ Stephen Lambrechts

Verdict

The best gaming mouse from one of the best companies making this gear.





KEYBOARD

\$279 | WWW.STEELSERIES.COM

SteelSeries Apex 7

Does this mark the apex of the trajectory of mechanical keyboards? Er, no.

SteelSeries's Apex 7 sits in a difficult position. It's an Apex Pro in all but its switches – and given that the Apex Pro's customisable hall-effect switches are (in theory) pretty much the most interesting keyboard controls of a generation, it's not immediately clear where the Apex 7, a linear-switched board that's a mere \$40 cheaper, stands. Is it good enough, without that gimmick, to sit so high on the price ladder? In its own bubble, at least, it's certainly a feast for the eyes. The RGB lighting hits you first, and it hits hard – both the key illumination and background coronae are absolutely superb, with a bright glow that seems to have been tuned away from the lurid side of RGB and into a more hip, warm, retro zone.

You can also map a custom logo on to the OLED screen at the Apex 7's top-right corner. That's neat, although using the clicky roller (which is lovely) and multifunction button (just OK) to interact with the

on-board controls only allows you to take a broad-strokes approach to your keyboard's configuration, with the rest done in software. SteelSeries does offer a number of apps to add a little more functionality to the screen – a Tidal “now playing” integration, and a Discord one, to name the two most notable – but these are pretty limited; it is at least a lot more useful when pulling up on-board profiles on machines that don't have SteelSeries Engine installed.

There's basically no wasted space, with that roller sitting on the corner of a very minimalist soft-edged aluminum and plastic base, and a magnet-attached wrist rest with a pleasing (though we suspect not wear-resistant) soft-touch coating. Even the hard-wired two-port pass-through cable, fully routable through channels on the base, is mercifully thin compared to the majority on the market, and the illumination around

that Type-A port is a wonderful touch. Additional design points go to the ABS keycaps and their bold, clear lettering, although the top-left offset marking of the main cluster may steer your fingers wrong if you're not a touch typist.

The eagle-eyed among our readers may notice that we've so far avoided talking about the switches in detail. Really, we almost couldn't bring ourselves to do so. Made in collaboration with Gateron, and clearly designed to ape Cherry's MX Red line, they're no doubt effective linear switches. They're fast and easy to actuate, they're very quiet indeed, and they're reasonably resistant. Problem is, they're also a little swampy, squidgy, and indefinite, with a soft (yet somehow slightly scrappy) feel that we simply do not like. We prefer Cherry's own Reds – and given that we find those to vary, depending on how they're applied, between passable and awful, that's quite the statement. Yes, switch taste

is a personal thing. You may lay one finger on this and fall in love. Who are we to say you're wrong?

However, there's one more aspect that's not so divisive: the price. Does the gimmick of an OLED screen, which is certainly much more useful in configuring the actuation point of the Apex Pro than it is messing about with lighting or sparse app integrations here, really put this keyboard in a higher bracket? No. It's special, solid, handsome, and very well designed, but it's way more expensive than it should be. Chop that screen away and slice off something like 40 dollars, and the Apex 7 would hit the mark. ■ Alex Cox

Verdict

Fabulous looks and excellent lighting options, but the switches aren't a win for us.





CPU COOLER

#299 | THERMALTAKE.COM

LABS BENCHMARK RESULTS

TEMPERATURES OVER AMBIENT (DEGREES C) - CPU CORE MAX FANS

THERMALTAKE FLOE RIING 360 RGB	42	
ALPHACOOL EISBAER 360 CPU	43	
EVGA CLC 280	37	
NZXT KRAKEN X61	42	

TEMPERATURES OVER AMBIENT (DEGREES C) - CPU CORE 50% FANS

THERMALTAKE FLOE RIING 360 RGB	48	
ALPHACOOL EISBAER 360 CPU	48	
EVGA CLC 280	43	
NZXT KRAKEN X61	51	

TEMPERATURES OVER AMBIENT (DEGREES C) - CPU PWM MAX FANS

THERMALTAKE FLOE RIING 360 RGB	18	
ALPHACOOL EISBAER 360 CPU	14	
EVGA CLC 280	7	
NZXT KRAKEN X61	14	

TEMPERATURES OVER AMBIENT (DEGREES C) - CPU PWM 50% FANS

THERMALTAKE FLOE RIING 360 RGB	21	
ALPHACOOL EISBAER 360 CPU	13	
EVGA CLC 280	10	
NZXT KRAKEN X61	18	

TESTBED: INTEL CORE I7-5930K @ 4.2GHZ, 1.20V.
PRIME95 V27.9, AVX FFT LENGTH 8K

Thermaltake Floe Riing RGB 360 TT Premium Edition CPU Cooler

Good cooling, quiet operation and a blaze of technicolour.

Henceforth abbreviated as Floe Riing 360 TTPE, Thermaltake's big RGB cooler certainly gives us a lot to talk about. It's RGB, and if you don't like that you can even turn the feature off. You could even leave the controller out of the circuit if you never plan to use its RGB features.

The box includes the radiator, three special nine-pin fans, a five-device controller hub, mounting brackets for Intel square-ILM interfaces and AMD rectangular-spaced mounts (both legacy and AM4), standoffs to support all those brackets, a lock ring to hold the brackets onto the head unit, fan and case mounting screws with washers, and adapter cables. The controller hub supports up to five devices, including RGB control for both the pump and fans, and fan control. A set of switches allows users to set a device ID from 1 to 16 on

the controller hub.

The head unit appears to be a standard Asetek part, right down to the smoothly lathe-cut CPU interface and factory-applied thermal paste. Tabs on the mounting bracket fit up through slots on its sides, and turning the bracket aligns its tabs to corresponding tabs on the base. A plastic ring can then be pushed up through some of those same slots to lock the bracket, if desired.

Installation begins by screwing standoffs either into the threaded brackets of your LGA 20xx or AMD motherboard, or the 115x/1366 backplate included with the Floe Riing 360 TTPE. The head unit is then placed with the ends of its mounting bracket over the standoffs, and secured with thumb nuts.

The finished installation looks great, with controllable colour patterns as well as the full RGB range of solid colours available.

In spite of its small pump, the Floe Riing 360 TTPE cools the CPU a little better than the Eisbaer 360. That's likely a limitation of the water block, since the Eisbaer is designed to support additional components. On the other hand, the Floe Riing 360 TTPE's temperatures are matched by the NZXT Kraken X61 and beaten by the CLC 280!

The tachometer shows us how the Floe Riing 360 TTPE ended up in the middle of the temperature metric: Its fans are really slow. Yet that could have a really good impact on noise, right?

We said that if you only wanted the most powerful cooler to fit your 2x 140mm mount (noise be damned),

the EVGA CLC 280 appears to be your best bet. Yet even its low speed noise level is greater than the maximum produced in the Floe Riing 360 TTPE. Thus, if you want the quietest possible cooling for your big overclocked CPU, the Floe Riing 360 TTPE appears to be your best bet. The Floe Riing 360 TTPE tops out our Acoustic Efficiency chart with the best cooling-to-noise ratio of any liquid system we've tested.

And if you're an RGB and overclocking and low-noise enthusiast, TT has you covered.

■ Thomas Soderstrom

Verdict

This provides the ultimate combination of RGB configurability, good cooling, and low noise, at a high price





KEYBOARD

\$399 | LOGITECH.COM

Logitech G915

Almost the perfect wireless keyboard.

It's hard to imagine a better wireless gaming keyboard than the Logitech G915 (\$250). It's a sleek, attractive peripheral with satisfying mechanical switches, cool low-profile key caps, a full media bar, two kinds of wireless connectivity and generous battery life. Moreover, it's not a tenkeyless model, it doesn't require AA batteries and it doesn't require tedious keyboard shortcuts to record macros.

Aside from a missing wrist rest (and really, a keyboard this thin and this expensive should come with one), there's very little to criticise in the G915. It's absolutely everything I wanted to see in a wireless gaming keyboard, no compromises required. However, \$400 is a lot of money, even for such a luxurious model, so weigh your options accordingly if you think you can live without a wireless setup.

The keyboard itself is gorgeous, with a brushed-metal-style finish over a sturdy plastic chassis. Each

key has an ultra thin key cap over a small mechanical switch, much like the fantastic Roccat Vulcan. There's a column of macro keys on the far left as well as three macro-profile keys and a macro record button above the F1 to F4 keys. Above the F4 to F8 keys, you'll find buttons for USB dongle and Bluetooth connectivity, as well as Game Mode (which disables certain buttons during play) and RGB lighting brightness. And, off on the upper right, there's a full media key section, complete with a volume dial. Every keyboard needs one of these, and the buttons work beautifully on the G915.

The G915 offers a choice of three flavours of a brand-new kind of switch – GL Linear, GL Tactile and GL Clicky. These new, low-profile mechanical switches aim to give just as much key travel as Logitech's Romer-G switches, but in a much smaller, more stylish

package. Each switch features a key travel of 2.7 millimetres, an actuation distance of 1.5 mm and an actuation force of 50 gf, but the feel of each switch is a little bit different. The GL Linear is quiet and requires you to press it all the way down to register commands, the GL Tactile is quiet and actuates before you hit the bottom and the GL Clicky is noisy and actuates before you hit the bottom.

You can connect the device via a USB dongle (which works flawlessly) or Bluetooth (which also works flawlessly – depending on how well your devices handle Bluetooth connections). Theoretically, this means you could have it hooked up to a PC and a smartphone simultaneously, or whatever other device combination strikes your fancy. Switching between them is as simple as pressing a button, and both connections work from dozens of feet away.

Then there's the battery life. Logitech claims that the device will run for up to 12 days of play (eight hours per day) with RGB features turned on or up to 135 days with lighting turned off.

As a wireless mechanical gaming keyboard, the G915 nails almost every feature perfectly. If you're serious about an uncluttered gaming setup without sacrificing any functionality, this is the peripheral you'll want as a centrepiece. ■ Marshall Honorof

Verdict

Pick up the Logitech G915 if you want a top-of-the-line wireless gaming keyboard with no compromises.





"If you didn't see the Moto G7 and G7 Plus' price tags before picking them up, you'd be forgiven for thinking they were far more expensive than the usual budget offering."

MOBILE PHONES

\$349 / \$449 (G7 PLUS) | MOTOROLA.COM.AU

Moto G7 and Moto G7 Plus

The best budget handsets you can get right now.

For many, the Moto G name has become synonymous with quality on a budget. Motorola has a knack for delivering relatively higher spec handsets than its competition at the same price point, and its new Moto G7 series is no exception.

Taking a lot of what made last year's Moto G6 series so great for users on a budget, the new Moto G7 range boasts a lists of refinements that make each phone feel even more like an upper mid-range product.

If you didn't see the Moto G7 and G7 Plus' price tags before picking them up, you'd be forgiven for thinking they were far more expensive than the usual budget offering. This is due to each device's premium feel and sophisticated design, offering sturdy metal frames sandwiched between glass on the front and back of the phones. Admittedly, there are a couple of spots where the

handsets cut corners to keep the cost down, but it's made those cuts in the right areas in order to minimise the impact on the user experience. They've each got big, bright 6.24-inch displays with minimal bezels and teardrop notches, too.

For the G7 Plus, the notch sports a 12MP selfie camera, while on the rear its got a dual lens 16MP + 5MP camera which allows for depth of field shots and a new 'cutout' mode which allows you to separate foreground objects from shots and applying your own zany backgrounds.

As for the regular G7, its cameras are a slight downgrade from the Plus model, with a 12MP primary sensor, 5MP secondary sensor and 8MP selfie snapper. However, we've found its photographic capability to be quite impressive regardless – you probably still won't find many smartphones with a better camera at this price point.

Though the entry-level

Moto G7 arrives alongside a slightly better sibling in the Moto G7 Plus, the standard handset is still an all-round terrific phone. It packs a quad-core Snapdragon 632 chipset (as opposed to the octa-core Snapdragon 636 processor found in the Plus model), but you wouldn't know it based on the snappy performance it delivers in everyday use.

Admittedly, the Moto G7's 3,000mAh battery is a bit smaller than some of the other handsets at its price point, such as the Nokia 7.1 with its huge 4,000mAh offering, making its battery life just 'fine' instead of 'great', with around 8-10 hours of juice with heavy usage. Of course, that's alleviated by the G7 Plus' 27W TurboPower charger, which will recharge the phone's battery by 40% in just 15 minutes.

The G7's is the same size and delivers more or less the same experience, however, while fast charging is available, it's

not quite as fast as the G7 Plus, offering a 15W charger in the box instead of its larger sibling's 27W TurboPower charger.

An improvement over its predecessor in every way, the Moto G7 Plus is highly recommended for those who want to balance functionality with affordability, while the standard G7 is good enough for most users who are just after a very affordable handset that offers nice build quality, great performance and decent photographic capability.

■ Stephen Lambrechts

Verdict

Superb value matched with good performance and a good camera.

G7



G7 Plus





"In addition to a larger 10-inch screen, the Nest Hub Max also has a camera that can be used not just to make video calls, but doubles as a security camera."

SMART DISPLAY

\$349 | STORE.GOOGLE.COM

Google Nest Hub Max

An excellent Echo Show alternative.

Later to the smart display market than Amazon, Google has taken a half-hearted approach to the category. Until recently, it's let others take the lead with Google Assistant-powered smart displays. Perhaps frustrated by the fact that only Lenovo, JBL, and LG have done so, Google released the 7-inch Nest Hub (\$299) last year, and is now following up with the larger Nest Hub Max.

In addition to a larger 10-inch screen, the Nest Hub Max also has a camera that can be used not just to make video calls, but doubles as a security camera. It also comes with a few new tricks, including facial recognition and the ability to track you as you move around the room. Reach to the back side and you'll find a switch that disables the camera and microphone. Flick it, and a small light to the right of the camera will turn amber, giving you a visual indicator that no one can watch or listen in.

With two 18mm, 10W tweeters and one 75mm, 30W woofer, the Max can put out a lot of sound for its size. Even at half volume, it easily filled my kitchen with quality audio. There was plenty of booming bass – almost too much at times – but high and midrange vocals were clear and sharp.

One of the advantages Google has over Amazon is that far more people use Google for far more things – such as storing photos, creating calendar invites, and looking up directions. And, once you connect your Google profile to the Nest Hub Max, all of that information is at your fingertips.

Setting up Face Match during setup was as simple as turning my head side-to-side for a few seconds in the Google Home app. After that, whenever I walked into the kitchen, the Nest Hub Max's camera would identify me, and greet me

with a cherry "Hello" or "Good Morning" on-screen. It would then show me missed messages, and commute times. The one thing that doesn't change is the background photos. They're fixed to the account of the person who set up the device.

Face Match, which can be set to recognise up to six people, is an opt-in service, so if you don't like the idea of your electronics knowing who you are, you don't have to set it up. Google also said that all of the face-recognition is done only on the device.

If you want to use the Nest Hub Max to make video calls, you're stuck using Google Duo. Yes, it's yet another communications app from Google, which your friends and family will have to install on their phones if they want to chat with you.

At \$349, the Nest Hub Max is about \$200 more than the Nest Hub, but

with that you get a larger display, more powerful speakers, and a camera which enables a host of other features. While I wish you could make video calls using services other than Google Duo, the Nest Hub Max's motion-tracking and facial recognition features are put to good use, and help separate it from other smart displays. I also appreciated the camera privacy switch, and I imagine others will, too.

It's taken Google longer than Amazon to come out with a 10-inch smart display, but the Echo Show now has a more than worthy competitor in the Nest Hub Max. ■ Mike Prospero

Verdict

A sleek design, smart motion tracking and handy privacy switch make the Nest Hub Max a great smart display.




software

» APPS FOR ALL YOUR PLATFORMS



Windows SOFTWARE



RadarScope

Superb weather information.

US\$30 | RADARSCOPE.APP

Extreme weather can pose a risk to health and also affect any plans that are outdoors, especially with hurricane season in full effect. If you're trying to stay up to date, then you might want a more detailed set of information than you get in most weather apps. RadarScope is a powerful weather tracking tool that lets you see an incredible amount of detail on your PC.

Before you scoff at the high price tag, remember that this isn't your average weather app. You won't see a weekend forecast with a beautiful Live Tile. Instead, what you get is incredibly detailed weather information. For example, you can view NEXRAD Level 3 data or Level 2 Super-resolution radar data. You can track the eyewall of a hurricane, see detailed tornado information, and track other detailed information from NOAA, AllisonHouse (if you have a subscription), and more sources.

You can use the app to keep track of extreme weather around supported areas and also save favourite locations.

RadarScope is aimed at people who are passionate about tracking weather information, whether they're using it for professional use or as a hobby. The interface of the app will take some getting used to for people who haven't used an app like this before, but the amount of information you can track is incredible.

One thing that you'll notice right away is that RadarScope doesn't smooth out radar information. You're getting raw data that you can interpret and digest yourself. I'm by no means a meteorologist, but it's fun to browse through the different information and notifications that the app offers.

While the Windows 10 version of RadarScope is powerful, it's been a bit behind its Android and iOS counterparts at times. Both the pro tier one and pro tier two were unavailable in the Windows version for a long time after the app's release on Windows 10. Luckily, these are available as options now. With these, you get dual-pane displays, hail contours, and cross-platform support among many other things. The tiers do require a paid annual subscription, but having the option is nice.

RadarScope has gained the pro features of its iOS and Android siblings, but it's unfortunate that the development of the Windows version of the app seems to be behind.

RadarScope isn't the type of app I'd recommend for everyone. It's a niche app that provides an incredible set of data to people who care about heavy-duty meteorological data. To those types of users, the relatively high price tag for RadarScope is worth it.

Sean Endicott



Mac » APPS



Adobe Lightroom CC

Now with more tools for editing and cataloguing images.

\$14.29/MONTH | ADOBE.COM | NEEDS MACOS 10.12 OR HIGHER

Back in 2017, Adobe rewrote its professional photo editing and cataloguing software, Lightroom, to bring feature parity to its desktop and mobile versions. In doing so they dropped a lot of features and functionality users had enjoyed in the much more mature Lightroom Classic CC (which is still on sale), but the developer has been steadily reintroducing features via updates since then.

The reason this is important to

mention is that Lightroom Classic CC is still a more 'professional' app – and still gets updates – but will invariably be discontinued, while Lightroom CC (the new version) will live on.

Priced at \$14.29 a month, Lightroom CC is both slick and streamlined, while providing more powerful and fine-detailed editing tools and photo management than Apple's own offerings in the Photos app.


On its release, Lightroom CC required you to store your original images in the cloud, but there's now the option to store them locally (though you do need to be online periodically for authorisation). Images you add to the Lightroom mobile app do get uploaded to your cloud space (which Adobe charges you to upgrade), but then all your images, tags, and edits do sync seamlessly between devices.

SOLID OPTIONS


Importing and tagging images is easy, as is creating albums and searching your library. The editing tools have been upgraded significantly since the

original release, and though still not quite as comprehensive as those in the Classic version, provide solid options. There's the usual stuff like exposure, white balance, tone curves, and selective colour editing, as well as detail and now texture – new in the 2.3 update – for enhancing the vivid qualities of an image.

Many camera and lens models are supported here with profiles for correcting distortion and there are preset looks provided – you can save your own too. Viewing and comparison options are good, and you can export images, though processing and format options are quite limited. For regular users, Lightroom CC is great, though professionals will likely want to stick with Classic CC for as long as they can. ■ Hollin Jones



Some of the more advanced tools include adjustment brushes, which let you paint changes onto selected areas.



App Store » iOS APPS

Calory


A simple way to track calories.

FREE (PREMIUM FROM \$23.49) | CALORY.APP
NEEDS IOS 11 OR LATER

Calory keeps track of your calorie intake (and despite its name it can measure in kilojoules). It uses age, height and weight to calculate a daily target based on your weight loss or gain goals, and you then enter your intake into the app to monitor your progress.

You can do that on iPad, iPhone or Apple Watch; the app supports Siri Shortcuts and can send and receive data to and from Apple Health. It provides a range of feedback tools including daily, weekly, monthly and annual reports.

The basic app is free, but the developer has since introduced premium features – such as barcode scanning for food logging, a large database of branded foods, a Today widget and custom reminders – that require a subscription from \$23.49/year. That's cheaper than MyFitnessPal, but that service



offers more detailed reporting, fitness tracker and fitness app integration. It also puts Calory head to head with other big hitters such

as Calorie Counter+, Lose It! and Cron-o-meter. In that company Calory feels a little, ahem, lightweight. **Carrie Marshall**

Grocery 2.0

Never forget your shopping list


FREE (PREMIUM US\$7.99) | CONRADSTOLL.COM
NEEDS IOS 11 OR LATER

This grocery-dedicated list app now has a Recipes section that allows you to create recipes and store ingredients. Recipes don't have to be recipes – you could create a group of items for your weekly shop, and add them to a list in one go. There's a timer to help you keep track as you cook, and an Apple Watch app so you can check off items on a list.

Grocery is closely tied to Reminders. If you allow access, when you create a Store or add an item to a list in Grocery, it syncs your list with Reminders. The app also uses Reminders' geofencing tools, so you can set the location of a shop in Grocery, and be reminded of your list when you get there. Lists are shop-based, and Grocery remembers

what you add to a list and uses that to make suggestions. If you prefer to keep apps separate, and don't allow Grocery access to Reminders, you won't get the syncing or geofencing benefits.

We found Grocery tricky to navigate. Adding and checking off items is fine, but adding a new shop is clumsy. If the idea of a dedicated grocery list app appeals, give Grocery a try but, for many, Reminders itself is fine. **Kenny Hemphill**





Google Play » ANDROID

Fimo – Analog Camera

Point 'n click.

FREE | N/A



A lot of apps tap into our collective nostalgia for the aesthetics of – let's be honest – lesser technologies, the static and hiss and bleed that we used to put up with and now sort of miss. Fimo is a digital recreation of an analogue camera, which transforms your screen into, curiously, a recreation of an early-2000s digital model that takes film. You can switch between eight different film stock types, and you'll have dust and scratches, frame shakes, light leaks, and other old-skool issues to contend with. At first it feels a little silly, taking shots on a modern phone with the field of view shrunk down and a big digital button that can't recreate the tactile

feeling of an old camera's little click-and-snap. But then you'll look at the shots, with their weird colours and date marks and artificial imperfections, and think 'oh, that's quite lovely, actually'. It's all pretty hipster-y, and not as cool as real film, but there's something, unexpectedly compelling about taking shots on digital film.



Moonmoons AR

Moon room.

FREE | [HTTP://WWW.ANNAMEREDITH.COM/ MOONMOONS](http://www.annameredith.com/moonmoons)



Moonmoons AR, a collaboration between composer Anna Meredith and designer Arthur Carabott, is essentially marketing for Meredith's new EP (in fact, a lot of AR apps are essentially marketing, but hey, it's a smart gimmick). Her new single, moonmoons, is presented here in its separate parts – you can place digitalised clay figurines representing the bass, soft synths, percussion, cello, arpeggios and textures around the room you're in through AR, and then walk around that room and experience the soundscape. Walk closer to the brain-like bass figurine and you'll hear it louder; put the arm representing a cello far away from the other elements to soften it. It's an experimental app, and with a good pair of headphones it's pretty impressive how well it works. It's enough to make you wonder why *The Flaming Lips* hasn't put out an AR app for *Zaireeka*, their album split across four CDs that you listen to simultaneously, yet. Honestly, the track is not my favourite (although Meredith has done great work in the past – her *Eighth Grade* soundtrack is excellent), but it's a cool app.



3D Geeks

Out of print.

FREE | [HTTPS://WWW.3DGEEKS.APP/](https://www.3dgeeks.app/)



This is a browser app for Thingiverse, a depository of 3D printer models where people show off what they've made and often share the blueprints. It's much easier to navigate on a phone than the mobile version of the website, with easy access to categories and tags, and no need to scroll side-to-side, only up and down. Preview images tend to load in at a low resolution, and the bland UI design makes the app a bit ugly, but it's still a very cool repository of what people are building and designing with 3D printers. You can load up 3D models of each component of many of the objects and examine them on your phone, rotating them to inspect every angle, and the licensing rights are listed on every object too. If you have access to your own 3D printer, this app can be very handy – it's fun for a casual scroll, but it's also useful if you need to show off a model to someone and you're not at your desktop.

■ James O'Connor



Game Boy Cartridge Holder Tray

AaronGraves



47



Collect



Share



Report

I needed something to store my Game Boy games rather than just sitting around. This holds 18 Game Boy games with a little room to flip through them or 19 completely full. I'm sure Game Boy color games would fit, but I don't have any to show

Sponsor

Hide >





Whonix 15

If routing all traffic through Tor doesn't satisfy your need for privacy, Mayank Sharma tests a resource-intensive solution with a novel approach.

FREE | WWW.WHONIX.ORG


Whonix is functionally very different from other privacy-focused distros. Here's a whistle-stop tour if you haven't heard of it: The distro ships as a pair of Debian-based virtual appliances that you must run simultaneously. There's the Whonix-Gateway that exists to configure and route all internet traffic via Tor. This machine has two virtual network interfaces: one connected to the internet via NAT, which is used to communicate with Tor, and the other connected to a virtual network.

The other machine is the Whonix-Workstation, which is the desktop you are supposed to use for browsing and other tasks. Firewall rules on the Whonix-Workstation force it to only connect to the virtual network and redirect all traffic to the Whonix-Gateway. This scheme prevents web apps from ever knowing the user's real IP address or accessing any information on the physical hardware.

The latest v15 release introduces several new features, including a couple that make Whonix easier to deploy. Starting with this release, the distro is available as a single combined appliance. When you import this, it automatically unfolds into the two virtual machines we mentioned earlier. By default, the VMs are configured to use four processors plus 768MB/2GB RAM for the gateway/workstation VMs. We reduced the processors to two for both VMs and increased the RAM to 1GB/4GB, and they still performed admirably well.

THOROUGHbred FOR PRIVACY

The other change is the introduction of a 'Live' boot option. When selected, the distro makes all writes to the memory instead of the VMs' hard disks, which means they won't persist after a reboot. Working from a Live environment adds a layer of security in addition to the distro's privacy features. However, you shouldn't rely on this alone to thwart data retrievals, since that requires additional steps, as the project's documentation notes.



Whonix 15 is also rebased on the new Debian 10 Buster release and has switched from the resource-intensive KDE desktop to the lightweight Xfce, which is a welcome change. We weren't enthusiastic about wasting resources running KDE on two VMs, and the switch to Xfce has made Whonix compatible with a lot more machines. Talking of underpowered machines, Whonix 15 has a CLI version that boots to a command line. You can theoretically use it to pull in another desktop environment, but the developers don't recommend this as it might pull in services whose privacy implications haven't been tested by Whonix.

On the Whonix desktop you get a cherry-picked collection of privacy apps. There's the KeePassXC password manager, zuluCrypt for encryption, qTox instant messenger along with some Tor-based tools such as the Tor Browser and OnionShare for sharing files. Highlights for us include the Firetools app to sandbox apps, the MAT2 metadata anonymisation toolkit, and the Kloak keystroke anonymisation utility.

In case a bunch of privacy tools don't convince you of Whonix's resolve to help you remain anonymous, perhaps

the distro's behind-the-scenes work to curb privacy leaks will. For instance, they have blacklisted uncommon network protocols out of the box to minimise chances of abuse. They've also taken steps to harden the kernel and have even sandboxed systemd services in a bid to limit the damage of a compromised service.

The best thing about Whonix is that its privacy hardening doesn't come at the cost of usability. The use of VMs eliminates the need for installation of any sort, which makes it very easy to deploy. Furthermore, the distro also ships with a handful of wrapper scripts to ease regular administration tasks. ■

thelab » long term memory

LONG TERM MEMORY

DESKTOP MEMORY IS BECOMING MORE AFFORDABLE;
CHRISTIAN GUYTON EXAMINES WHAT THAT COULD MEAN FOR YOU.



RAM is a core component of any computer. With DDR4 now well and truly entrenched in the memory market and DDR5 on the horizon, manufacturers are looking out in different directions: massive 64GB and 128GB sets running across eight channels, and faster speeds than ever seen before. Frequencies only continue to rise, with companies such as HyperX and Corsair pushing commercially available sticks at well over 4,000MHz.

But it isn't cheap. RAM price-per-GB skyrocketed between 2016 and 2018, although recent trends have shown that DDR4 kits have slowly lowered in price from the latter half of 2018. RAM offering 3,000MHz and above remains expensive, with a seemingly endless list of brands and product lines to choose from. It's daunting.

The vast majority of modern memory comes packaged with integrated heatsinks, with many sporting multicolored LEDs to boot. Some beefier packs even come with their own coolers, mounting awkwardly above the DIMM sockets. With an open mind and an even more open wallet, it's easy to spend a silly amount of money.

Choosing RAM in this day and age can be a laborious process, but fear not – as always, APC has got your back. Over the next few pages, we're going to break down the technology and nuances behind random-access memory, and discuss which kits are best for which builds. A discerning eye is important when buying RAM; what you pay for is not always what you get.




Size does matter

When looking at memory kits, there are three key elements to consider: capacity, speed, and latency. We're going to take a look at capacity first, as it tends to be the biggest deciding factor when choosing your RAM. Industry regulation from standardisation body JEDEC (more on that later) means that memory only comes in certain multiples: 2GB, 4GB, 8GB, and so on, doubling each time.

For consumer use, the standard maximum capacity of DDR4 chips is 1GB, due to limitations in transistor size, although a theoretical 2GB per chip is possible, so the largest commonly available capacity per DIMM slot is 32GB. 64GB

sticks are available, but not for general consumer use; they're more likely to be found in server farms. With a motherboard supporting full memory capacity, the current maximum available for desktop systems is 128GB.

Of course, 128GB is far more than the average PC will ever need. RAM should be chosen based on the requirements of the system; it can be one of the more expensive components of a fresh build, and it's easy to overspend on capacity that isn't required. With that in mind, here's our guide to how much RAM your system should need.



4GB

■ Consider this a bare minimum. If you're building on a seriously tight budget, or just looking for a basic PC for office work (think word processing, not giant spreadsheets), 4GB of basic DDR4 will do. As memory frequency has less of an impact on Intel processors than AMD ones, we recommend opting for slower RAM, such as DDR4-2133, paired with a cheap Intel Core CPU. This memory won't cost much – a stick of 4GB RAM shouldn't run you much more than \$20 right now. Dual-module kits are also an option, using two 2GB sticks to improve performance, but these can cost more, and the difference is likely to be minimal. Rigs such as this also make for affordable home theater PCs or cryptocurrency mining units, although some cryptos can be CPU-intensive, so a slightly better processor is a good idea.

8GB

■ If you're looking for a standard PC, 8GB of RAM is the way to go. Capable of handling any office work and low-end gaming, 8GB of dual-channel memory will serve any basic build well. It's worth bearing in mind that if your GPU doesn't have much VRAM, some games attempt to cache large texture files to the system's DRAM. If you're running an older graphics card with less than 4GB of video memory, 16GB of memory could be a better call. This only applies for more graphically demanding games, though; if your poison is lightweight indie gaming, 8GB is plenty. Older triple-A titles ought to be fine, too, if you've got an appropriate GPU to back you up. Make sure you go for dual-channel; a single stick of 8GB DDR4 will work, but the performance difference will be more noticeable in games.

16GB


■ For mid-range gaming supported by a decent GPU, 16GB of DDR4 is a necessity. Games are becoming increasingly demanding when it comes to RAM; if you're looking to play in anything higher than 1080p, 16GB may not always be enough. The uptick in multiplatform games being ported to PC has resulted in a lot of poorly optimised triple-A titles hitting our screens, too, which have a tendency to gobble up RAM. 16GB should be able to cope with other tasks, though; image manipulation and intermediate video-editing work shouldn't be an issue. You'll likely have two options: 4x 4GB or 2x 8GB. Doing your research is worthwhile; one kit might appear better value for its performance, but if your CPU and motherboard are only going to support dual-channel memory, you won't be able to reap the benefits of a quad-channel kit.

Stepping Back Through Time

■ The concept of random-access memory has been around since the 1940s, although it didn't exist in its current form until 1968. When we talk about RAM here, we mean DRAM – dynamic RAM – the kind of memory you're probably familiar with: Sticks that plug into your motherboard and use binary capacitors and transistors to temporarily store readable data. Early RAM worked differently, using electrically charged spots on cathode ray tubes or networks of magnetised metal rings. The core principles remain the same, though – the points of positive or negative data could be read in any order, therefore denoting them as "random access."

The current model of RAM, using transistors to store data as combinations of charges, came into commercial use in 1970 with the Intel 1103 chip. This memory requires a constant supply of electricity, with the minuscule charges on the thousands of transistors needing to be refreshed hundreds of times per second. Unlike storage memory (such as that found on traditional SSDs and HDDs), RAM is volatile; once the computer is powered off, the charges dissipate, and any data is lost. While non-volatile RAM does exist, it tends to be too slow for normal memory needs, and can suffer from other issues.

Modern DRAM uses solid-state memory on silicon circuitry, unlike the ceramic and magnetic components of the Intel 1103. It functions in much the same way, though, losing data once the power is turned off. Nowadays, RAM is used for a wide variety of computing functions, most commonly for graphical tasks, be that rendering 3D objects in real time in a game, or editing video footage. DRAM is found in more than just the heatsink-equipped silicon sticks you can buy in varying sizes and kits; it operates inside modern graphics cards, games consoles, even calculators.




Say hello, Grandpa... This is the Intel 1103 chip, the precursor to modern DRAM as we know it.

32GB

■ Some RAM kits have got significantly cheaper over the past year. In early 2018, 32GB of DDR4 could set you back \$500 or more, but prices have dropped by around half since then. At the moment, a kit of four 8GB sticks should only set you back \$100-\$200, even at higher speeds, such as 3,600MHz. With this new affordability in mind, 32GB (or more) becomes a more distinct option. If you're looking to edit 4K video, use 3D rendering programs, or play high-end games at maximum graphical settings, you'll want at least this much memory to play with. For gaming purposes, 32GB will be overkill for most modern games, but also serves to future-proof your system. Again, if you're rocking four 8GB modules, as opposed to two 16GB ones, make sure that you've got support for quad-channel memory.

64GB+

■ It wasn't that long ago that the notion of putting 64GB in a desktop PC was faintly ridiculous, in terms of both cost and necessity. Today, though, 64GB of high-speed RAM (over 3,000MHz) can cost as little as \$250, making it feasible in a high-end system for gaming or editing tasks. While 32GB modules are available, we recommend a quad-channel kit of four 16GB sticks; dual-channel runs the risk of RAM bottlenecks. 64GB will give any system serious oomph, but if you have money to burn, 128GB is an option too. It's hard to see where this would be required – advanced high-res videography, perhaps – but the option exists. DDR5 promises to up that maximum to 256GB, which will likely only serve a niche role for heavy data analysis tasks and video processing in 8K and beyond.



■ We know, memory speed can be complicated. Speed and frequency are often treated as interchangeable statistics, because they are directly tied to each other; RAM, like many other components, operates at a specific frequency. This frequency represents the number of times per second an individual transistor gate can open and close. With memory operating at, say, 1,000MHz, these transistors can cycle at a fairly nippy 1,000 million times per second.

If you're the memory-savvy type, though, you might be more familiar with the MT/s figure than MHz. Shorthand for megatransfers per second, MT/s was introduced as a metric for memory speed when DDR (double data rate) was first produced, because the newer format of memory was able to process data at twice the speed by registering an additional data transfer on the rise and fall of the clock. Frustratingly, many manufacturers chose to ignore this standard, meaning that current-generation memory quotes frequency figures that are double their true operating frequency. If you examine memory speed using system status analysis software, such as CPU-Z, you'll find that your 3,000MHz memory is actually running at around 1,500MHz. MT/s is technically the more accurate figure, but MHz is more widely used.

Obviously, more speed is better. Commercially available memory speeds have been rising steadily over the last few years, with the current DDR4 standard reaching as high as

4,800MHz (or MT/s, to be more precise). Higher speeds reduce your real-world memory latency, and enable the transferral of more data within a set period of time. This enables superior performance in rendering functions and other tasks, although not all processors are affected equally; if you're rocking an Intel CPU, you'll find that increasing your memory speed will have a far smaller impact than on an AMD processor. For now, then, capacity remains more important than frequency.

Memory speed does impact another important feature of your RAM, however. The channel bandwidth of any memory kit is determined by memory speed; for those not in the know, the bandwidth is the maximum amount of data that can be transferred between the memory and CPU within a set period of time. This is calculated by multiplying the transfer speed (in T/s), the system memory bus (in bits), and the number of channels the memory is capable of accessing simultaneously. So, for a set of dual-channel DDR4-3000 memory running with a CPU using a 64-bit bus, the channel bandwidth would be 3,000,000,000 x 64 x 2, which is 384,000,000,000 bits per second. A gigabyte per second equals to 8 billion bits per second, so we end up with the far more palatable figure of 48GB/s. This is the absolute maximum data rate that this system could achieve; attempts to do anything demanding faster transfers than this would result in memory bottlenecks.

Assuming that you're not going to

meddle too much with your CPU, that means that data transfer is determined by more than just memory speed. When choosing what RAM to put in your system, consider the advantages of a quad-channel kit over dual-channel (provided your CPU supports it of course – most mainstream CPUs are resolutely dual-channel). To be fair, dual-channel memory support represents a good standard for the average system, but quad-channel support effectively double your channel bandwidth, so if you're using your system for tasks that demand massive data transfers (such as 4K+ video editing, or large-scale texture production), four can be the golden number.

Channel support is tied to your motherboard, too. Many high-end boards support quad-channel memory, but if you're considering a memory upgrade for a slightly older system, do bear in mind that your motherboard might only support dual-channel. Four DIMM slots is no guarantee of quad-channel support; many four-slot boards have two primary slots (which logically should be A1 and B1, but are sometimes A1 and A2 or A2 and B2, depending on the manufacturer), which provide full dual-channel speeds. This works both ways: Plug in only one stick of dual-channel memory and you'll get half the channel bandwidth. Smaller motherboards, such as the ITX format, have only two slots with dual-channel support, due to their layout specifications, even on X299.

Quality control

■ The Joint Electron Device Engineering Council (JEDEC) is the body responsible for maintaining industry standards when it comes to numerous solid-state components, including RAM. With over 300 members, including Microsoft, Intel, and Apple, JEDEC pools resources to push hardware development forward.


Its most important contribution to RAM standards is the adoption of a unified connection format for motherboards (the DIMM slot) to ensure compatibility, but that's not all. Standards for data transfer rates and latencies are determined by JEDEC, so consumers choose between speeds such as DDR4-3000 and DDR4-3200, instead of 3060, 3111, 3187, and so on.

It's simple enough in principle, but these constraints ensure memory is easily put into straightforward brackets for our convenience.

JEDEC's members are also responsible for producing new hardware within the RAM universe. The up and coming memory standard is DDR5, which was ratified in March 2017 (not to be confused with GDDR5, a memory standard for GPUs). Some manufacturers (such as SK Hynix) have confirmed that their DDR5 memory is complete, and desktop DRAM modules should be available later this year. This next-gen memory will be able to reach higher speeds than any units from DDR4, and draw less power per module, thanks to

integrated voltage regulators.

While it might sound like a minor improvement, this shift in standards has implications beyond RAM design. With voltage controlled on the memory sticks, motherboards will no longer require voltage regulators for RAM, which will enable manufacturers to focus on other features, but could cause compatibility issues in the not-too-distant future. DDR5 will double the bandwidth available, so 128GB may become more mainstream, and 256GB builds will become a possibility for high-end tweakers. Memory speeds will hit a staggering new maximum at 6,400MHz, the fastest available within DDR5.



JEDEC doesn't actually require heatsinks to be integrated on RAM.

Latency & timings



Problems? Chances are it's your RAM. Try manually setting the frequency and timings in the BIOS.

■ Memory speed isn't the only thing that affects your RAM. Latency and memory timings encompass a huge variety of key statistics that affect memory performance, and many of those numbers will be meaningless unless you're a die-hard hardware-head. The one to look out for is CAS (column address strobe) latency, which is an indicator of how quickly the RAM stick responds to a memory call.

Bear with us, because it's going to get mathematical. CAS latency by itself isn't necessarily a helpful figure – in order to get a genuinely useful number, there's a formula you need to use. Real-world latency = (2,000/RAM speed) x CAS latency, with speed measured in MT/s and latency measured in nanoseconds. So, if you've got a memory kit that runs at 3,000MT/s and has a CAS latency of 16ns, the kit's

real latency is 10.6ns.

However, faster memory tends to come with greater CAS latencies, so the actual performance gains can be minimal, even when you seek out higher frequencies. The ideal memory kit has a high frequency and a low CAS latency, and overclocking is an option, but there's a lot of potential messiness there. If you're running an Intel system, variation in speed isn't going to make much difference; real-world latency in DDR4 and beyond is starting to become increasingly homogenised, regardless of frequency, so buying high-speed RAM can become a waste of money. The impact on some tasks is more notable; gaming and 3D rendering is less affected by memory speed than large-scale data analysis.

CPU differences



■ We've mentioned that memory specs have different impacts on Intel and AMD processors. In broad terms, memory speed hasn't historically made much difference to processor performance, but with AMD's Ryzen series of CPUs, that norm changed in a big way. Due to conventional processor architecture, Intel CPUs (currently the Core series) only see minimal gains from memory speed; a couple of percent at best in the vast majority of benchmark functions.

With Ryzen, though, AMD did something different: The interface that connects the processor cores together – called the Infinity Fabric, if you subscribe to funky hardware marketing jargon – is controlled from the integrated memory controller on the chip. What this means is that when supported by faster RAM, the memory


controller can run faster, and transfer data between the cores at higher speeds, improving overall processor performance.

This was a stumbling block for Ryzen in the early days, because when the processor line first became available in 2017, memory was significantly more expensive. The effects of memory speed on Ryzen chips works both ways: Slower RAM can result in poor performance compared to similar Intel Core CPUs. As memory has become cheaper, though, Team Red looks more and more attractive. A few hundred additional MHz can make a difference in processor performance of 10 percent or more, which can be pumped up even higher when combined with CPU overclocking to produce some very impressive numbers.

The bottleneck

■ In the early years of commercial RAM, before the turn of the century, there was a serious disparity between CPUs and memory. Referred to as the "memory wall," a problem arose where CPUs were developing faster than RAM technology, with RAM speed only increasing by 10 percent in the 14 years leading up to 2000, while CPU speeds increased by approximately 55 percent. For a long time, it looked like memory speed limitations might cause a significant bottleneck in the improvement of computer performance.

Since then, CPU speed increases have slowed, due in part to physical design barriers around heat dissipation and die size in new chip designs. With the improvement of caching tech to speed up key processes and bridge the gap between CPU and RAM speed, it looks like the "memory wall" might be a fallacy.



What else can my RAM do?

■ RAM does more than simply boost your PC's processing power. If you're particularly software-savvy, you may be familiar with the concept of a RAM disk: a block of RAM that is used by your system as a temporary secondary storage drive. Using specific software (such as DataRAM RAMDisk), it's possible to partition a segment of memory to be used as a drive. These virtual drives can be incredibly fast – far faster than conventional SSDs or HDDs – and theoretically have no limits on read/write times.

The main drawback of virtual RAM disks is that they lose all stored data when powered down, making them particularly vulnerable in the event of system crashes or power failure. They're also an expensive way of producing storage space to work on

files – RAM is simply more costly than physical drives. However, there are other advantages. RAM disks can be ideal for working with decrypted versions of encrypted files, and you can set up regular backups of your data to a physical drive to get around the risk of data loss on crashes. It's possible to keep data stored on a persistent storage drive, set to recreate it on the RAM disk when the system boots. Most interestingly, RAM drives can be used to access compressed files without having to decompress them first, using systems such as cRAMfs on Linux. With memory capacity ever increasing, RAM disks may become more relevant in the not-too-distant future.

Other mechanics of RAM have been slowly integrated into modern

operating systems for ease of use. Virtual memory, for example, involves the partitioning of hard drive space as a "paging file," which is then used for storing excess data when the physical RAM is operating at capacity. Windows does this automatically, using the paging file as additional memory space to prevent program crashes. This function can be disabled in Windows 10 to free up disk space, although it's not recommended if you're running anything less than 32GB, and does put you at greater risk of system instability. Modern BIOSes also incorporate "shadow RAM," where ROM functions instead use DRAM locations to speed up access times. The benefits of doing this manually are minimal, though, particularly with modern ROM chips.



Best RAM choices for your system

Whatever platform you're looking to build, memory is going to be an integral part of it. We're going to break down the different potential purposes of new system builds, and assign the best RAM for the job in each scenario. RAM might have become cheaper over the last year, but it's still all too easy to splurge when you don't really need to.

Bear in mind that these are only loose recommendations; follow the specs as a guideline, rather than the manufacturer. As always, though, we advise you to stick with recognised brands, such as Corsair, Crucial, G.Skill, and HyperX. If you can't find the manufacturer's name listed anywhere in this magazine, tread lightly.



HIGH-END WORKSTATION

**HyperX Predator
64GB (4x 16GB)
@ 3,000MHz \$530**

■ To be honest, there are no wrong answers here. If you're hankering after a PC for seriously demanding digital creative work – whether that's large-scale 3D modeling or 4K video editing – the sky's the limit when it comes to memory capacity. We've gone for a mid-range 64GB kit here, with a real-world latency of just 10ns, but if you've got the cash

to splash, consider faster or even larger modules. Corsair's Dominator Platinum series provides some excellent 64GB and 128GB kits, for example. Make sure your CPU and motherboard support quad-channel memory bandwidth, as is the case with any four-module RAM kit.




OFFICE WORK

**Corsair Vengeance
LPX 8GB (2x 4GB)
@ 2,400MHz \$70**

■ If you're just keeping it simple, this is all you need – 8GB of RAM will comfortably handle light office work, be it spreadsheets or slideshows, and is more than capable of supporting your usual array of web-browsing activities. We would be a little wary of loading up on Chrome tabs, however, with only

eight gigs behind you, but this should be able to carry you through everyday tasks. We contemplated going for Team's slightly faster 2,666MHz model in the T-Force Vulcan line instead, but the price difference was minimal, and for baseline performance, 2,400MHz is more than enough.



BASIC WORKSTATION

**Team T-Force
Vulcan 32GB (2x 16GB)
@ 3,000MHz \$1,300**

■ Even with low to mid-level photo and video-editing work, 32GB of memory is the way to go right now. It won't necessarily be cheap, but the affordability of higher-capacity DDR4 kits has improved no end over the past year, perhaps due to the looming release of DDR5. This much memory will run rendering tasks well,

and equip your system for high-definition videography. If you're looking for a fast system to run tools such as Photoshop or After Effects, look no further. 2,400MHz will be plenty for Intel systems; if you're on an AMD machine, it may be worth splashing out for a little more.



BASIC AMD GAMING

**Ballistix Tactical
16GB (4x 4GB)
@ 3,000MHz \$150**

■ The Ballistix Tactical memory isn't particularly exciting to look at (perhaps check out Ballistix Tactical Tracer RGB if you're looking for some more eye-candy-worthy modules), but it provides excellent performance for a sensible price. With a high frequency and a CAS latency of 16, this kit will pair nicely with a mid-range GPU for

some straightforward 1080p or even 1440p gaming. It won't necessarily be able to hit maximum settings in every game you play, but well-optimised titles will give you no problems. Should you decide to upgrade your GPU, a memory upgrade is worthwhile, too, and the Tactical or Elite series are good ones to stick with.




HIGH-END AMD GAMING

**Corsair Dominator
Platinum RGB 32GB (4x
8GB) @ 3,600MHz \$600**

■ This is where it starts to get expensive. We're cranking up the speed to make sure your AMD system's latency stays low, and paying the extra price for 32GB ensures your build is future-proofed. Most games don't use 32GB – 16GB is usually recommended – but with GPU improvements, that could change. With this much capacity, there won't be any

danger of running low on available memory, even when you're bashing your system with ultra 4K graphics. The Dominator Platinum RGB is a lovely looking kit, too, with fully addressable LEDs on every module.




BASIC INTEL GAMING

**Ballistix Sport LT 16GB
(2x 8GB) @ 2,400MHz
\$120**

■ There's nothing much to write home about here, but this Ballistix kit will do the job and do it well. Memory speed is less of a factor when you're looking at RAM for Intel systems, so you can save yourself a few dollars by opting for a slightly more conservative frequency. The Ballistix Sport series is Crucial's

more affordable memory, but that doesn't make it any less effective. It also runs at just 1.2 volts, meaning it has some of the lowest power draws found within the DDR4 standard. It also has a CAS latency of 16, which means it keeps your real memory latency at 13.3ns.



HIGH-END INTEL GAMING

**G.Skill Trident Z
RGB 32GB (2x 16GB)
@ 3,000MHz \$530**

■ Once again, we don't need to push to boat out too much when it comes to memory speed in an Intel machine. Unlike the Dominator Platinum RGB kit, this set comes with two 16GB modules instead of 4x 8GB ones, so if you're running a system with more than two DIMM slots, make sure you're using the prioritised

connectors for maximum speed. This is G.Skill's own RGB offering, and they're pretty sweet to look at. A CAS latency of 16 means this might not be the fastest RAM around, but it's an extremely good price for 32GB. The extra capacity can come in handy for streaming memory-intensive games.

SUBSCRIBE & GET APC FROM ONLY \$5.79 AN ISSUE!

12 ISSUES FOR ONLY \$79!

WHY SUBSCRIBE?

- Take advantage of our trial subscription offer of six issues for \$49 if you choose to auto-renew*
- Get 12 issues of Australia's favourite tech mag for only \$79!
- Save even more and get 24 issues for \$139 – that's just \$5.79 per issue!
- Free delivery of every issue straight to your door
- Free full-version software downloads with every issue

SAVE UP TO \$99.80!



**FREE DELIVERY TO YOUR HOME
OR OFFICE EACH MONTH**

HURRY

OFFER ENDS NOV 4TH 2019!

TWO EASY WAYS TO SUBSCRIBE TO APC MAGAZINE

VISIT WWW.TECHMAGS.COM.AU

OR CALL (02) 8227 6486

TERMS & CONDITIONS: This offer expires November 4th 2019, and is available to subscribers within Australia only. Please allow 6-8 weeks for delivery of your first issue. For full terms and conditions, see: <http://techmags.com.au/style/images/FutureTC.pdf> Please see our privacy policy on page 3. *Six issues for \$49, then \$49 charged to your credit or debit card every six issues thereafter on auto-renewal.

PERFECT PASSWORD PRIMER

Human laziness, complacency, general failings, and unending ingenuity to subvert systems have undermined passwords forever. Neil Mohr shows that it's time to next-gen your security, and save all your passwords!

In any system, humans will find a way to subvert the rules for their own gain. It's a recurring theme: from skipping class, cheating in games, fixing the banking system, and exploiting welfare, down to stealing from work, and gambling in Las Vegas. So it shouldn't be a surprise that most humans choose the lazy option when it comes to password security.

It's too easy just to say humans are lazy, and pin terrible password security on that. Humans will always attempt to game systems at both ends, which is why passwords are being hacked and people want to avoid memorising them in the first place. It's not good enough to just use a

strong password; we now live in a world where so much is tied up with our online lives that hackers are doing more and more to break into online accounts.

So, we're going to look at the whole subject of password security, and hence online account security, how we can choose better passwords, how we can make our lives easier but also safer, how we can increase protection beyond passwords, and how hackers attack and use your account details.

Some of this advice you'll have probably heard before, some of it is hopefully new, but by the end, you'll be terrified enough that you'll overcome your human propensity for laziness, and get up off your bottom, and protect your passwords!




Let's start by scaring the heck out of you. There has been a host of research on password security, made much easier by the huge number of database breaches in recent years. One of the widest-ranging research efforts was by Google, published mid-2017 in its 'Data breaches, phishing, or malware?' Understanding the risks of stolen credentials' paper (<https://research.google.com/pubs/pub46437.html>).

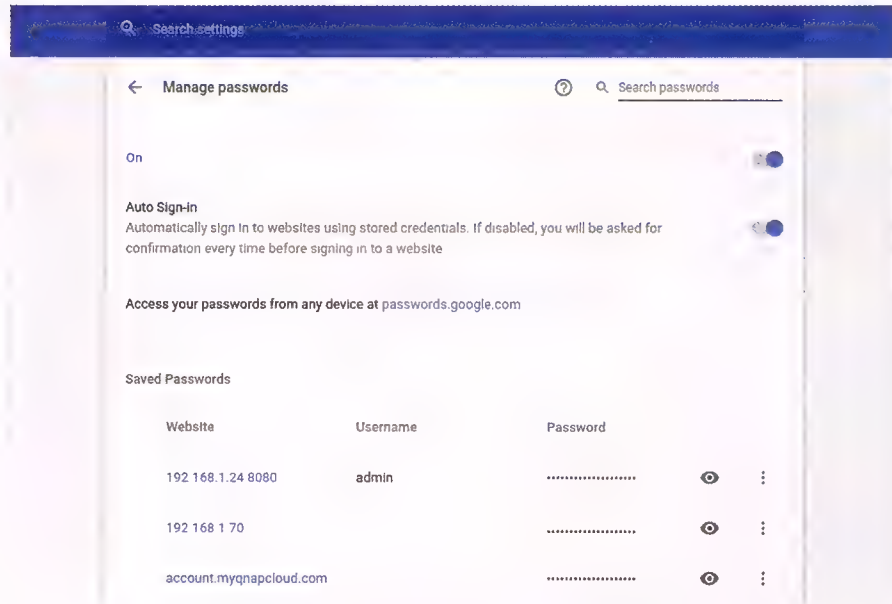
Some sweet nuggets from that little cache include the fact that the researchers swept over 1.9 billion non-unique account usernames and password credentials. These were taken from a variety of leak sources, but mostly blackhat security forums. An insane 76 percent of these could be (or already were) easily converted to plain text (best practice is to hash and salt stored credentials). The researchers were able to reverse an impressive 36 percent of hashed passwords, using a suitable keyword attack.

Of those stolen via phishing attacks, 49 percent were from Americans. When it came to pure security leaks, 39 percent of credentials were linked to Americans. As for credentials stolen via keyloggers, Brazil accounted for the highest amount in this group at 18 percent.

Credential leaks appear to happen regularly these days. The largest ever remains the initially hushed-up Yahoo leak, back in 2013, in which pretty much every Yahoo account ever was leaked, totaling three billion. Adult



A well-integrated mobile app handles online accounts and local app logins.



Google Chrome can manage and synchronise passwords across your devices.

Friend Finder topped 412 million, the 2017 Equifax breach hit 143 million users, while MySpace, Adobe, LinkedIn, Dropbox, LastFM, NexusMods, and more have all had their leaks, each individually releasing millions of users' details.

These leaks show the basic issue with password reuse: If you are using the same password across all of your online accounts, it just takes a single leak to contain that password, and all your accounts are compromised.

SCIENCE SAVE US!

So the logical next question is: What makes a secure password? Entropy. That means the level of randomness in a system. By randomness, we mean true randomness. It's not enough to have a long password, it has to be long and truly random. We stumbled across a lovely analysis of poor password creation by WP Engine, at <https://wpengine.com/unmasked>.

Among many things, it shows that password creation is a truly gender neutral pursuit – men and women can live equally in a world with terrible passwords. It's worth a read – as well as being funny and interesting, it highlights a lot of general failings us humans have when it comes to creating passwords.

The obvious beginnings are the standard bad passwords: '123456', and any variation thereof, 'password', 'qwerty', and again any variation of those. Obvious stuff. Another issue is that many password guidelines are bad – forcing upper-case letters, for instance, because people tend to just capitalise the first letter; forcing at least one number, because people just put '1' at the end; or demanding a password of at least a certain length, because people use a pattern on the


keyboard. For example, the seemingly random 'ADGJMPTW' is actually people typing 2-9 on a smartphone number pad. This is what happens; it doesn't make for strong passwords, because they're easier to guess with a much lower entropy.

It's that human nature thing. One group wants to protect a system, another wants an easy life, so it circumvents the rules, a third group is attempting to break in. So, how do we make stronger passwords?

A few years ago, a Dropbox engineer, Dan Wheeler, wrote a blog post (<https://blogs.dropbox.com/tech/2012/04/zxcvbn-realistic-password-strength-estimation>) referencing Randall Munroe's classic XKCD 'correcthorsebatterystaple' cartoon (<https://xkcd.com/936>) on password entropy, alongside Mark Burnett's earlier research (<https://xato.net/perilous-password-policies-3ad038aa33b1>). Dan's post is long and complex, and focuses on what makes a strong (high entropy) password, alongside how that can be easily measured and communicated to the user. Conversely, Randall's point was part password strength, but he was largely extolling the need to make password systems human-friendly.

We don't get into the nitty-gritty too much, but effectively a good password is a long password. Simple. The XKCD approach is to use a memorable series of words or phrase, usually painting an odd scene, such as the classic example of a 'correct horse battery staple', but it could be a phrase such as 'basketball on a unicycle flying high'. From a memorable point of view, this taps into techniques used by memory experts. Of course, throw in some numerals and symbols, and it's even stronger.

None of this addresses the core issue, though, that we now have a seemingly



Most password managers work through a browser plugin, as that's where we all live now.

endless number of online services to use and try, each with their own password and login. Even making passwords memorable wouldn't solve the innate issue of having to remember them all. The solution to that is a password manager.

MICROSOFT SAVE US!

The obvious place to start is to ask what tools Microsoft provides in Windows 10 or beyond. The blasé answer is none. Windows doesn't come with any specific password management tool. What Microsoft has done is build a basic password manager into its Edge web browser, which enables those saved passwords to be synchronised across your Windows accounts, and to any other instances of Edge you might be running – Edge on Android, for example. But what if you're not on your Windows account? What if you're not on Windows at work or home? What if you're using Chrome or Firefox? What

if you want to create a strong password? What if a site is demanding non-standard credentials? The Edge password manager is basic at best, and fails to help in all but the most simple of scenarios.

BROWSERS SAVE US!

How about a browser that you might actually want to use? The most widely used web browser at the moment is Google Chrome, and it comes with a reasonable password manager built in. To kick off, if you've logged in with your Google account, it syncs passwords across all of your devices and operating systems (Windows,

macOS, Linux and Chrome OS), and is, of course, available on Android and iOS devices, too. Google also offers an online service at <http://passwords.google.com> that enables you to log in and access passwords from within any web browser on any device.

Beyond passwords, Google Chrome also supports intelligent form filling. You can enter a default address, phone number, and email address, and if Chrome thinks you're filling those details in, it offers to try to auto-fill entire forms or individual elements. Alongside this it has a secure credit card storage option, if you find the idea of Google having your credit card details a good one... What could possibly go wrong?

Chrome's basic but smart enough password manager means that as long as you're logged in to your Google account, a right-click in a password field offers the option to generate a secure password. There are no options, it just provides a random password of suitable length. It then, of course, saves and manages that login and password for you. You can set exceptions and manage the saved passwords from within the Advanced Settings. Access to these is protected via a system prompt that ask you to enter your regular Windows password.

The other main browser worth a mention is Mozilla Firefox, although it doesn't provide much more functionality than Microsoft Edge. It synchronises your password across its

Alternative managers

We're going to suggest three: the commercial password managers – the ones that crop up most often, and generally offer similar features at similar costs. So, it's down to you which you feel works the best. The first is 1Password.com; prices start at US\$2.99 per month for a single user, and US\$4.99 for a family account for four people.


Dashlane (www.dashlane.com) is a major player and offers a number of plans that will suit a wide range of consumer and business uses. The basic plan is free

for a1, but is limited to working on a single device. The paid home plan is US\$3.33 a month, and supports unlimited devices, password sharing, and Yubikey 2FA support. A business plan starts at US\$4 per user, and includes remote management, group sharing of passwords, and enhanced 2FA options, among others.

Our final option is Keepersecurity.com; it has a similar structure to Dashlane, but opts to have yearly subscription pricing. The basic single-user package is US\$29.99, effectively US\$2.50 per month, and offers

unlimited devices with sync, unlimited password storage, secure cloud backup, fingerprint access, and web access with support. A family option costs US\$99.99 for five users.

There are, of course, a host more password managers. You can find a vaguely independent list on Wikipedia at https://en.wikipedia.org/wiki/List_of_password_managers, or our sister website TechRadar maintains a list of free (ish) password managers at www.techradar.com/1325845.



feature » perfect password primer

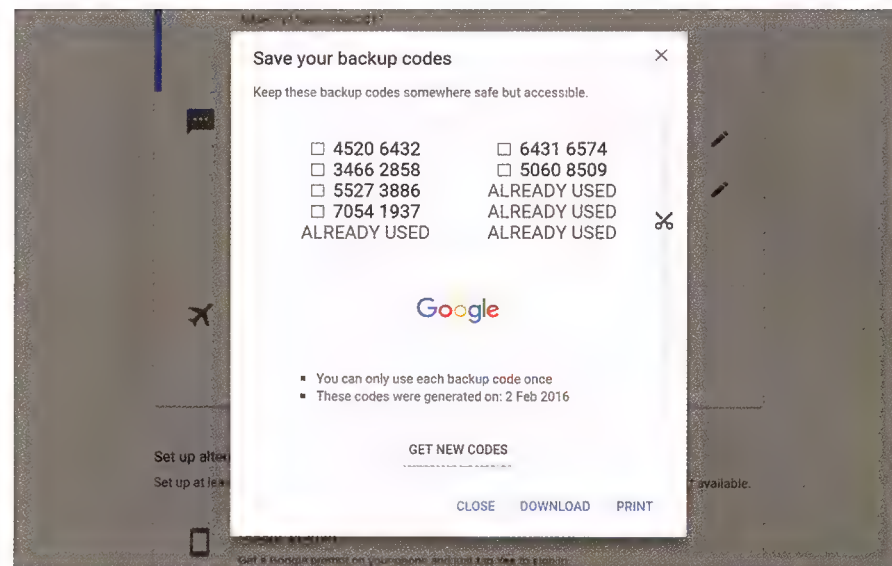
own Firefox account to any supported OS or devices. It doesn't support forms, and won't generate passwords for you, but it does offer an independent master password to protect access.

LASTPASS SAVE US!

So, browsers generally leave us wanting more when it comes to real password management. You could get by with Google Chrome, but we can do better. Enter the world of dedicated password managers. We're going to concentrate on LastPass (www.lastpass.com/) in this main section – for alternatives, see the various boxes throughout the feature. It's certainly one of the prime options to look at when it comes picking a password manager; the free account provides the majority of the functions people require, while the paid-for premium edition is just \$4.22 per month.

LastPass is a multi-device, multi-OS, multi-browser tool. It largely works as a browser plug-in for the big five: Chrome, Firefox, Safari, Edge/Internet Explorer, and Opera. It runs on all the main desktop OSes: Windows, macOS, Linux, and OpenBSD. And for mobile devices, there's Android, iOS, and Windows Phone support. Plus, you can log in to your LastPass password vault through any browser, too.

The name LastPass comes from the fact that your LastPass account password will be the last password you need to remember; it'll take care of all the rest. So, what does it do? The main



One-time use access codes enable emergency access if you lose your 2FA device.

use is via a browser plug-in and apps on your mobile devices. As you create new accounts or log in to existing ones, LastPass stores your usernames and passwords, or offers to generate new passwords as accounts are created.

The LastPass password generator is pretty sweet. You can choose any length, up to 100 characters, select any mix of upper and lower-case, choose the quantity of numbers to use, and decide whether or not it should use symbols. There are also advanced options for avoiding ambiguous symbols or attempting to make your

new password pronounceable.

Accounts can be grouped into categories and searched. You can add notes and opt for automatic logins, to help speed you into accounts. LastPass also supports automatic filling of standard forms: names, addresses, emails, phone numbers, bank accounts, credit cards, and custom fields. There's a general encrypted notes tool, too. It also offers a warning for insecure forms, plus a security audit of existing accounts that you might import from your web browser. Two-factor authentication is supported, too – we cover this later. The

The open-source way

We like open source, it might not always produce the most polished looking software – (cough! GIMP! cough!) – but for security software, its publicly scrutinisable code tends to provide peace of mind, while largely ensuring a project can continue even if the original organisation or developer packs up their bags, as the source code stays open to be picked up by someone else.


We're going to highlight two strong open-source candidates. The first is the well-known KeePass.info. Don't confuse this with the largely defunct **KeePassX.org** that hasn't been updated for two years. There is a fork called **KeePassXC.org**, which is up to date, but we'll stick with KeePass as it has many unofficial reskins for Android and other devices that make it easier to use.

KeePass isn't as straightforward to use as the commercial alternatives, the main trip you need to jump through is for multi-device synchronisation, where a

third-party cloud service, such as Dropbox or Google Drive, needs to be roped in. But the breadth of support, custom options, and feature set is unmatched, with no fees to pay.

The other strong open-source option is Bitwarden (<https://bitwarden.com>). It offers an excellent free model, a US\$1-a-month family subscription that supports

up to five users, a basic business version at US\$5 per month, plus enterprise options. It's a modern, slick, open-source application supporting Windows, Linux, macOS, all the major browsers – including Opera, Tor and Brave – plus Android and iOS devices. If you're a fan of open source, we suggest you give it a look.



Online gaming service Steam has offered secondary authentication for many years.

paid-for service expands password filling to other applications, on both Windows and Android – LastPass can fill in a program's password and fields; handy these days as more programs require some type of login for

subscriptions.

Once installed, it's worth looking at LastPass's Preferences. At APC we like to add a log-out time, so if you leave your PC alone, it logs itself out after an hour, for example.

LastPass isn't the only password manager to offer all these features, but it does provide an idea of the expanded capabilities, improved security, and genuine ease of use a dedicated service offers.

But with great power, comes great responsibility. A service such as LastPass requires you to use a strong master password; if that is hacked, someone will gain access to every service you use. Or if you forget the password, LastPass can't restore your master password, as it is never sent it in the first place, as a deliberate security measure. It might sound like we're putting all our eggs in one basket, but what do you want – an individual basket for every egg, or just one awesomely secure basket?

Even security-conscious businesses can be hacked and run into issues (https://en.wikipedia.org/wiki/LastPass#Security_issues). LastPass has had a number of vulnerabilities exposed, and even had a successful hack attack on its own servers in 2016. But as a security conscious company should, LastPass offers bug bounties – which it honours – and is open about breaches with its users, while using best practice on its secure storage. So even the server breach in 2016 meant

users' own passwords remained secure.

GABEN SAVE US!

A question that security experts started to ask a while ago was: Is there anything better than passwords? You won't like the answer, but it's unlikely we'll ever be completely free from passwords, although two-factor authentication offers a respite. Called 2FA or TFA for short, the premise is that people require two forms of identification, AKA authentication; typically, the first is a strong password (something you know), and the second can be your fingerprint, your delightful face, or a generated one-time-use passcode (something you have).

As you might have spotted, many phones are now using fingerprints and face scans to authenticate access. Financial institutions are even happy enough that NFC-equipped phones can authenticate credit card transactions. Some systems (LastPass, VeraCrypt) accept USB flash drives and specialist USB keys (Yubikey, www.yubico.com) as a second form of authentication.

Steam is a good example of a service or organisation that was quick to provide two-factor authentication. You should be familiar with its email verification, which sends a code to your default email address, while more recently, Valve introduced the Steam Authenticator app for smartphones and tablets, which is required if you're messing around with its trades. ■

Authenticate with Google


We're going to highlight the Google Authenticator as a prime two-factor authentication (2FA) tool. Install it on your phone, then fire up the service you want to protect with 2FA. LastPass uses Google Authenticator (among others) to do so. Using that as an example, head to LastPass Vault > Account Settings > Multifactor options. Click the Edit button and Enable it. Use the Barcode option to display a QR code. On your phone, open Google Authenticator, tap + and click 'Scan barcode', then point your phone at the on-screen QR code. You should now find that LastPass has been added to Google Authenticator.

Anytime someone tries to access your LastPass Vault from a new device, they must enter the current code (this updates every 30 seconds) before access is granted. There's an option to allow access for 30 days, before another 2FA prompt appears on that device. Without your phone, there's no access.

2FA is great, but does have a couple of caveats. Firstly, the service using it

mustn't be monumentally stupid. It was discovered that some services could have 2FA circumvented by simply calling up and providing some basic personal information, making the use of 2FA more security theatre than actual security. The second is that potentially the 2FA code is locked to the one phone. Technically, a rooted Android device enables you to transfer the Google Authenticator secrets, but if you get a new phone, ensure you disable 2FA on your services, before wiping your old one. Otherwise, you could be locked out of vital services. LastPass keeps the QR code secret, so you can have it on multiple devices, and transfer it later on – provided that you can log in, of course.

Because of the magical nature of cryptography maths, many services also offer offline access options, usually via a bank of single-use passcodes that you can print or scribble down, and keep safe somewhere. This means that if your phone is lost, stolen or broken, you will still be able to access your accounts.





FACEBOOK'S

Chris Lloyd examines the social media giant's impact on the finance industry with Libra.

Facebook's Libra is nearing launch, and if the plans bear fruit, it's going to be the most disruptive thing to come out of Silicon Valley for years. Let's start with what Libra isn't. It's not a cryptocurrency, you won't be able to mine coins, and it's not an opportunity for speculation or investment. It does use similar blockchain tech to drive it (see "The Blockchain" box), but it's a

closed system designed for stability. It's also designed to be fast, something bitcoin struggles with (it can barely manage seven transactions a second; Libra will start at 1,000). It isn't a bank account either, something Facebook is keen to stress for legal reasons, so you won't be earning interest on your account, or get credit.

Libra is a stable private blockchain



NEW MONEY

digital currency, and will be accessible through Facebook applications. It is not unlike a PayPal account in operation, but without a link to your bank account. Facebook has developed a subsidiary, Calibra, to handle it. You will be able to store Libra "coins" (or whatever units of Libra are called; it's not clear yet) on the Calibra wallet application, which links into WhatsApp

and Messenger. Libra will also be available through Google Play. Once an account is verified using a government-issued ID, you will be able to spend hard cash to buy Libra, either electronically or by handing over actual folding stuff at the planned participating outlets. Facebook is in talks with chains of convenience stores to make cashing in and out readily accessible.

You can then send your Libra "coins" to another Libra wallet, or spend them to pay for goods and services online. Eventually, you'll be able to tap and pay with your phone for your coffee, Uber ride, or whatever you want. You can also cash in your coins, getting hard currency in return. It's not clear yet whether this will involve a fee, although it looks likely. Libra is not

meant to be a vehicle you put money into; it has been designed to be used as money itself. There are plans to make it possible to buy Libra at ATMs, too. Any fees will be low, certainly considerably lower than using traditional financial services, although we don't know exactly how low. The Libra website talks of "low or no" fees, and of potential rewards and bonuses for choosing to pay with Libra. As the Libra website says, "Libra's mission is to enable a simple global currency and financial infrastructure that empowers billions of people." That's big talk, but then Facebook does like to frame its projects in grandiose language.

When you buy a Libra coin, what you buy into is the Libra reserve fund, currently backed by assets worth around US\$1 billion. The reserve fund is held in major world currencies and US Treasury bonds, which are traditionally pretty stable. By spreading the risk, the value of a Libra coin should remain fairly flat. And because you are likely to buy and then spend your Libra coins quickly, any



Get used to these wavy lines; Libra introduces a new global currency.

long-term drift is relatively unimportant. Initially, a Libra coin will be worth about one US dollar.

Libra will be much more than hipsters tapping phones to pay for skinny whites. All you need is a smartphone, and you have access to banking-like services. That's something that nearly two billion people don't have. Libra can power the kind of micro-transactions that make all the difference in developing countries. The potential market is massive: There are 100 million people in Pakistan alone with no access to a bank account. The market for international remittance is a target,

too (see "Sending Money Home" box). Libra has potential to do a bit of good, as well as earn Facebook and friends a lot of money.

Libra may be revolutionary for the financial market, but the technology really isn't. The widespread use of Internet-enabled devices and the development of the blockchain ledger made it inevitable that Libra, or something similar, would appear eventually. Using money on the Internet used to be an awkward process – if you remember a time before PayPal and debit cards, you'll know. Libra is the logical next step, where money is converted into tracked,

"When you buy a Libra coin, what you buy into is the Libra reserve fund, currently backed by assets worth around US\$1 billion. The reserve fund is held in major world currencies and US Treasury bonds, which are traditionally pretty stable."

The Blockchain

What is the blockchain? It is tempting to say it's a mass of complicated math you don't need to understand.

Essentially, it's an electronic ledger that tracks the movements of digital tokens between wallets – apps that store the tokens. The tokens can represent anything: ownership of a share or proof of identity, for example. Each transaction is verified and added to the list, or chain. The wallets have two cryptographic keys, one private and one public. A transaction is encrypted using your private key and the recipient's public key. At the other end, the process is reversed, so it can only be unscrambled by the intended recipient, and they know it's from you. The scrambling involves that complicated math we mentioned. Each transaction becomes part of the whole, because it's mathematically linked to previous transactions, so it's virtually tamper-proof. It's a neat system, which can run on a distributed network (traditional cryptocurrencies) or private systems.

While Libra uses much of the same underlying technology as bitcoin, it works differently. Coins are created as required, and "burned" (destroyed) when they leave the system. This is quite unlike bitcoin, which has a limited number of coins that, in theory, never get burned, although an estimated quarter of all bitcoins have been irretrievably lost. Although it is called the

Libra Blockchain, it doesn't use blocks of data, but rather treats the data as a whole. This simplifies things, and speeds it up. It is also the reason why any transition to a distributed model will be very difficult, despite the idea that Libra will be opened out in the future. Facebook has termed Libra as a "decentralised programmable database," although this hardly helps make things clearer. It has been widely reported as having an initial transaction speed of 1,000 per second. If true, it looks a little inadequate, unless there is some batch-processing involved. Visa runs an average of 150 million transactions per day, which equates to around 1,700 a second. Given Libra's expected use for multiple micro-transactions, and the hoped-for size of the user base, it'll need to expand capacity at some point.

The banking world has been employing blockchain technology for years, but hidden away behind the scenes, managing business-to-business transfers, clearing, and settlement transactions between banks, for example. While the adventures of bitcoin, with its crazy valuations, has captured the headlines, back-room blockchain systems have been working their way into the servers of big business.





Look – you can send money while on the back of a bike. A cheesy, but slick, promo video stresses Libra's simplicity and global nature.

transferable data, with the blockchain keeping a record of who has how much. Your dollars are free to flash across the world, along with the endless text messages and selfies.

Cryptocurrencies were always a threat to Facebook, as well as a huge opportunity. Facebook's billions are made selling ads to influence what people buy. Any company that developed a successful stable cryptocurrency would threaten that market by building a valuable set of data that tracks spending habits – just what advertisers want. If that company had the reach and power of, say, Google, it could take a sizable bite out of Facebook's business. Libra means Facebook can keep transactions within a controllable suite of Facebook apps, giving even more reason for a business to buy advertisements there.

Most tech enterprises attract little attention from the world at large. Facebook carried on quietly fueling the social media boom for years before attracting attention from legislators. The fact that Libra is special is clear when, months before even a limited launch, it attracted comments from the financial chiefs of the G7 nations.

So, why all the fuss? Because money is power, and Libra side-steps so much regulation, and that has implications. If Libra gets as big as it plans, it will put a great deal of power into the hands of those who control it. Libra may sound libertarian – many aspects are – but it has the potential to be intrusive, and worse. The prime mover behind Libra being Facebook hasn't helped bring



The Calibra wallet app links into WhatsApp and Messenger.

positive reactions either. The company has a terrible record on privacy; only a few months ago, it was fined a record \$5 billion for passing on private data. The records generated by Libra will be kept by Calibra, and can't be shared with Facebook's servers without authorisation using an opt-in system. However, there are plans for a friend-finder search function, which does require the linking of Calibra transactions with Facebook accounts. The trick, it seems, is to get users to give permission to share information, through incitements of some sort. Collecting data is what Facebook does, and Libra will be part of that process.

Your financial records are a goldmine for advertisers. Can we trust Facebook not to exploit them? In surveys, Facebook does not score well at all. A recent one by NBC and the *Wall Street Journal* had 60 percent of people not trusting Facebook with personal data, more than twice the figure for Amazon. We may not trust Facebook, but we haven't stopped using it. That financial data isn't just a juicy target for business, either – security forces would be very interested, too.

The advent of Libra has drawn a great deal of attention from US Congress, too. Banking is a heavily regulated industry, while tech companies have generally received a light touch. That looks set to change. Facebook has moved into different waters here. In July, the Banking Committee held hearings and questioned the head of the Libra project, David Marcus, for two days. It wasn't the friendliest of meetings. The Senator for Ohio said that the company showed "breath-taking arrogance," and that Libra was "a recipe for more corporate power over markets and consumers." Facebook denies that it is a bank, and therefore is free from bank regulations. Which is just as well, because one committee member said, "If you're a bank, we (will) regulate the heck out of you."

There is not a lot of love for Facebook in the US government. It has, as yet, failed to give concrete answers on how it will implement anti-money laundering measures. Basing the organisation in Switzerland has drawn heavy fire, too (see "The Libra Organisation" box.) Mr. Marcus claims



The Libra Association

Facebook will not run Libra directly, but through its subsidiary Calibra, it will be a partner in the body that does: the Libra Association.

This is a non-profit, independent organisation dedicated to running, controlling, and promoting Libra, as well as managing the reserve fund. Facebook hasn't gone alone for a number of reasons, trust being one of them. As the project leader David Marcus admitted, "If we were controlling it, very few people would want to jump on and make it theirs." The association has 28 founding members, and among them are such big names as PayPal, MasterCard, Visa, eBay, Uber, and Spotify. There are also some names you won't know, unless you are familiar with blockchain development and venture capital. This is just the initial lineup – it is expected to grow to over a hundred members by the time Libra goes live.

Each founding member paid in at least \$10 million to get a seat at the table. Decisions on the running of Libra are made by voting, with no single member getting more than one vote or 1 percent of the total votes (whichever is the greater). The idea is that Libra cannot be dominated by a few of the larger members (in theory, anyway). Each founding member is responsible for running a validator node, which maintains the integrity of the blockchain. Members will also earn interest on cash held in the system – Libra coins, or tokens, created, but as yet unspent – which could become a useful amount if the system grows as hoped. It is also used to pay for running the blockchain: server time costs.

The Libra Association is based in Switzerland. Why? Because Switzerland is a major hub for international money transfers, and Swiss banking laws are unlike any other, and notorious for their secrecy. Swiss banks were also notorious for money laundering, particularly after World War II. Thankfully, the country has cleaned up considerably since, but is still seen as one of the safest places in the world to store money. It is estimated that a quarter of all global cross-border assets are held there. Libra is expected to be regulated by the Swiss financial regulator, although it is still in preliminary talks, and has yet to reply to the Swiss privacy regulator when asked to provide details on its operations.

the location was chosen to give Libra an "international place," not to skirt regulations. Facebook has also said it won't launch Libra before the regulatory issues have been resolved.


It's not just in America that Libra faces problems. The French financial minister, the governor of the Bank of England, and the Swiss Data Protection Commissioner have all issued statements that warn of the need for regulation. Many countries have laws banning or limiting the use of cryptocurrencies; one huge market Libra was hoping to hit early was India, where such laws have blocked its launch for now.

The potential for criminal use is another frequently raised concern. There are tight banking regulations to stop money laundering – what does Libra have? Promises – and rather vague ones so far. Although regulated by the Swiss government, it says it will comply with all US tax, money-laundering, and fraud regulations. These include registering with FinCEN (Financial Crimes Enforcement Network), which aims to stop terrorist funding and money laundering by analyzing data collected during financial transactions. It will also comply with the international anti-money-laundering regulations (AML). These include confirming people's real identities, and sending the identity of wallet-holders to the relevant government agencies. Facebook has also applied for an export license, and jumped through a number of other legal hoops. This hasn't stopped calls for the project to be stopped until every concern has been addressed, though.

Facebook claims that outdated financial regulations have restricted commerce – for example, the banning of the majority of banks to trade in cryptocurrencies. It says that Libra offers real innovation, freeing capital quickly and easily anywhere it is needed, rather than being stuck in the banking system, which can take days to make a simple transfer. Exactly how Libra will be regulated has yet to be resolved. Whether or not it is treated as a bank or as an exchange-traded fund will make a big difference. Facebook says Libra will complement existing financial systems, and it isn't a bank, because it does not interact directly with its users. A moot point, which will probably require lawyers to unravel.

We can expect some more toing and froing, but it will take more than a few Representatives publishing open letters calling for Facebook to halt the project to stop this train. However, this is the biggest ever company to get involved in cryptocurrency, and on such a large scale. We can expect an unprecedented level of scrutiny before launch, which will have ramifications right across the world, and on any future such project.

Cryptocurrencies are often subject to scams and hacks, which have cost coin exchanges millions.



The Libra wallet app, Calibra, in action. You can send money as easily as a text (almost, anyway).



The Libra Association is keen to stress its use for sending small amounts of money internationally.

Sending money home

A stated early target for users of Libra is the migrant worker who sends money home to support their family.

While the individual amounts may not be large, the scale of the whole business is huge. The World Bank estimates that US\$715 billion is moved around the world annually, supporting over 800 million people, and making major contributions to the economies of many nations. The established wire transfer services, such as Western Union, take a cut of this. The global average for transfer fees is 7 percent, which is quite a hit: Seven bucks for every hundred of hard-earned wages disappear when you simply send money home. There are currently few ways to avoid these fees, and they have a disproportional effect on wealth. UNESCO (the United Nations Educational, Scientific, and Cultural Organisation) has called for this to be reduced to a maximum of 3 percent to aid the developing world.

The infrastructure Libra needs is in already place – the use of smartphones is high, even in countries that lack other basic services. It's not unusual for more people to have smartphones than have access to a bank account. A quick bit of math also shows why this is such a tempting project for Libra. If Facebook manages to capture half of the market, and charges a more reasonable 2.5 percent, that equates to nearly \$9 billion a year. This is one reason why it had been planned to launch Libra in countries such as India early, although this has hit a regulatory roadblock for now. Offering free or low-fee money transfers is also an excellent way to get Facebook into emerging foreign markets.

Libra will not be immune; it uses open-source software and allows for third-party wallet applications, which can cause problems. Already, there are reports of dozens of fraudulent pages on Facebook and Instagram offering to sell Libra, often at a tempting discount. Facebook is taking them down, but it is a game it cannot ever really win.

It's difficult not to see the irony here – Libra involves real money, where trust is paramount, and the world of Facebook is awash with fakers of various hues. The very platform that Facebook uses to promote Libra is not trustworthy itself. About five percent of all active Facebook accounts are fake. Earlier this year, Facebook attempted a clear-out, and deleted a staggering 2.19 billion accounts. Although many were profiles and removed at the point of creation, it is still an alarming number. Facebook will have to get much better at

identifying and removing fake Libra pages and ads. It has been here before, too – there were so many scams involving cryptocurrencies that it banned all ads for them in January last year. It relaxed the rules a little this spring, allowing related products and services to be promoted, but you still can't sell cryptocurrencies or ICOs (initial coin offerings) on Facebook.

Can it all go horribly wrong? That is a frightening, if unnecessarily pessimistic, prospect. If it did start to wobble, who would bail it out? No government is under any obligation to do so. Some doom-mongers say Facebook is at risk of developing a system that is too big to fail, but is decentralised to the point where there is nobody obliged to catch it if it does. There is little need to worry unduly, though – Facebook will have done its research thoroughly, and the project is clearly well funded, and has a

powerful group of backers.

Libra is going to be big. It is also going to hurt the core business of many large international organisations, as such disruptive technology has a habit of doing. However, it is also almost inevitable. If not Facebook, then somebody else, and Libra is unlikely to have the market to itself for long either. Ready or not, such currencies are coming. Libra should be ready to launch next year, although exactly where and what service it will initially offer is unclear. In the meantime, there is going to be a good deal of sorting out, as the regulators and Facebook bargain over exactly how Libra is to be regulated and implemented. In a few years, we will probably be wondering what all the fuss was about. The basic idea is solid, practical, and genuinely useful: a stable world currency designed from the outset to be used in a digital world. ■

DEEP INTO STORAGE SPACE

At the edge of the final storage frontier, Anthony Howes looks at how the universe's greatest storage system is about to unfold.


We all know SSDs are coming to kill old spinning storage, but for many of us, a bank of HDDs still offers the capacity we need – 8K footage, a mass of multi-megabyte photos, and who knows what other encyclopedic-level storage demands the discerning APC reader has? – without totally breaking the bank.

There's no doubt SSDs are catching up with their clunky spinning brethren on capacities, but the cost of a 4TB cutting-edge SSD drive would buy you nearly 30TB of spinning storage. If only there were a way to combine the storage space of old spinning drives with super-fast solid storage.... Well, say hello to Microsoft Storage Spaces. A modern solution for an ancient problem: Where to store all your old junk without enraging the other half. Storage Spaces are the Microsoft

implementation of what the rest of the computing world generally calls storage pools.

It's an attempt to abstract away the messy workings of RAID, mirrors, and JBODs, and present to the user a simple unified pool of storage, with the option of adding caching drives, the ability to mix and match drives of varying capacities, the bonus of device redundancy, and an easy way to manage your spaces.

It seems like the perfect solution, so why isn't everyone using storage spaces? Well, Microsoft being Microsoft, it went and had to try to hobble a useful system for standard Windows 10 users. So, we're here not only to explain how to get the most from your storage spaces, but also put back the missing functionality, no matter what Microsoft build you're looking at.



feature » deep into storage space

Everyone's looking for ways to make their busy lives easier and simpler. Because laziness is the true mother of invention, and the world is full of sysadmins who don't want to be hassled by failing drives, there came a solution: Storage Spaces. These are a way of creating virtual drives from a storage pool of real hardware storage, no matter what form that hardware takes.

No longer do you need to have a mess of multiple drives. No longer do you have to mess with RAID systems. With Windows 10, you can just pile together a bunch of storage, and have it appear as a single storage space that comes with a bunch of additional friends with benefits.

Storage Spaces is included with the consumer release of Windows 10; the server product Storage Space Direct has a few additional features, but is only available on Windows Server 2012/16/19. By and large, the extra features are aimed at enterprise and data center type usage. One feature, however, is support for SSD tiered storage and write-back cache – essentially, it enables you to provision a dedicated SSD drive(s) to cache a larger Storage Space. It appears that Microsoft removed GUI-based support in the consumer version of Windows 10, but it is possible to use PowerShell to manually specify a drive. So, we'll look at this once we've covered the basics.

We're going to look at your options for configuring Storage Spaces; part of that examines how the different options work for combining physical drives together into what Microsoft calls a storage pool. The combinations can get pretty varied, especially when you factor in expanding pools, and how virtual Storage Spaces can be "provisioned" across them. We want to

manage pools, replacing drives, rebuilding them, and optimising when needed, while we also look at using an SSD to enhance write performance. And, of course, we can throw a few benchmarks at our pool to see how it performs.


THE SCIENCE STUFF

Before we get stuck in with provisioning a few drives, we have to do some theory on how pools are created and why they exist, because it'll save some head scratching, and having to rebuild your Spaces when you realise you've not done an optimal job. If you're hot to trot on RAID and the like, you might want to skip ahead. As we've alluded to, Storage Spaces came about from admin demand for simplifying – aka abstracting away complexity – storage management. It's not hard to imagine. You have a data

center full of server racks, packed with spinning drives that are continually failing. You not only want a system that can prevent data loss when drives fail, but also a system that simplifies managing all this storage as a single block, or at least a series of blocks.

Taking the first part of that problem, one of preventing data loss, you want to add redundancy to your storage: If one drive fails, you don't lose data, and everything can continue as if nothing had happened. Storage Spaces, and traditional computing, offers two options: drive mirroring and parity storage, aka RAID.

Mirroring is the easiest to understand. The basic configuration is when you have two identical drives, and the same data is written to both – if one drive fails, you have a spare copy on the other. Simple? Well, mirroring can get more complex: Imagine each




Ensure you clear any drives before you try to add them to a storage pool.

SAME SIZE, BUT DIFFERENT

Can you use drives of different sizes? Yes, but no. Clear as mud – just how we like it. The short and basic answer is that in any parity or mirrored pool, you can use any size of drives you like, but the maximum capacity of the pool is determined by the smallest drive.

Take a mirror consisting of a 1TB and 2TB drive – the maximum mirrored capacity is 1TB; the 2TB drive is limited to 1TB of storage. For a parity array, take a 4TB, 4TB, and a 6TB pool of drives. The overall capacity would be 12TB, because the 6TB is limited to the 4TB of the two smaller drives, though space available for storage would be 8TB, as a third is reserved for parity data, as we've explained.

If the occasion arises, you can more cleverly match your different-sized drives in mirrors to maximise your space. For example, if you had two 4TB and two 6TB drives, you'd match up one 4TB and one 6TB to create one side of a 10TB mirror, and the remaining two drives on the other.



You can assign multiple Storage Spaces, even if they're bigger than the physical storage.


side of the mirror as a storage "column." Storage Spaces enable you to combine multiple drives together in each storage column. One column mirrors the other, but with multiple drives, which is handy when you want to expand an existing mirror with more drives. Storage Spaces have also extended the idea of the basic mirror to a three-way mirror, so rather than just two mirror storage columns, you have three. The theory being the window of failure on a complex two-way mirror is big enough to warrant additional redundancy in the system.

The issue here is that a standard mirror wastes space – 50 percent of your total storage is lost as redundancy. A three-way mirror is even worse, with 66 percent "lost" to redundancy. Isn't there a more efficient option?

Enter parity storage, also known as RAID (Redundant Array of Inexpensive/Independent Drives).

Avoiding going into the actual math, parity uses a logic trick – it's used in error checking, too – that means that by comparing two bits being stored, you generate a third parity bit. These are stored over your base three "parity" drives – the minimum number required. Bit one is written to drive one, bit two to the second drive, and the parity bit to the third drive. If any one drive now fails, you can get your data back by reversing the parity creation process, called rebuilding.

From this basic system, you should hopefully see how you can build it up to writing two bytes of data and then a parity byte – the actual amount of data is variable and typically called a "stripe," as data blocks and parity blocks are striped across the array of parity drives. Four, five, or more drives can be added, and the data and parity data is simply spread across the drives. In this case, 33 percent of your total storage is used as redundancy, so it's more efficient than a mirror, especially



We found Microsoft parity-based Storage Spaces to be hot garbage, performance-wise.

SOLID STATES

Windows offers two features to accelerate Storage Space performance: write-back cache and SSD tiers. On the consumer release of Windows 10, these can only be implemented through PowerShell when the Space is originally created. We failed to add an SSD tier to a parity array of HDDs, but did have success with a single SSD and a simple array, plus a mirrored SSD and HDD array. We'll look at how to create both.

We're starting the process from a blank slate, so first remove your Spaces and pools. We'll be using PowerShell, so search for that, and run it in Administrator mode. Step one is to create a variable that stores the available blank drives; run this command without `$disks =` to see the output. We can then create our pool from that list.

```
$disks = Get-PhysicalDisk
-CanPool
$True
New-StoragePool
-PhysicalDisks $disks -
StorageSubSystemFriendlyName
"Windows Storage*"
-FriendlyName "Pool"
```

Next, we create two more variables that store the details for what will be our SSD tier and HDD tier for drives.

```
$ssd = New-StorageTier
-StoragePoolFriendlyName
"Pool" -FriendlyName SSDTier
-MediaType SSD
$hdd = New-StorageTier
-StoragePoolFriendlyName
"Pool" -FriendlyName HDDTier
-MediaType HDD
```

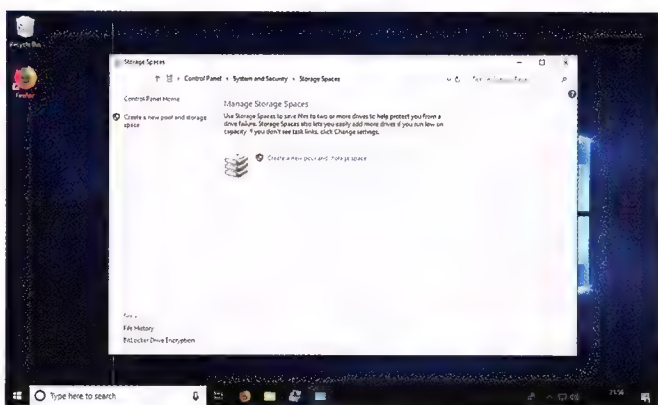
Finally, we create a Storage Space with the `New-VirtualDisk` command – you'll need to change the `-StorageTierSizes` values to suit your devices.

```
New-VirtualDisk
-StoragePoolFriendlyName
"Pool" -FriendlyName MySpace
-StorageTiers $ssd, $hdd
-StorageTierSizes 400GB,
3000GB -ResiliencySettingName
Simple -WriteCache
```


Open the Manage Spaces control panel and you'll see a new unformatted space. You can format and manage it within the GUI from here. If you have two SSDs, you're in luck, because you can create a cached mirror space. Repeat the last set of commands, but this time around, simply enable the `ResiliencySettingName` as `Mirror`. We assume that if you had three SSDs, you'd be able to enable parity.

```
New-VirtualDisk
-StoragePoolFriendlyName
"Pool" -FriendlyName MySpace
-StorageTiers $ssd, $hdd
-StorageTierSizes 400GB,
2000GB -ResiliencySettingName
Mirror -WriteCacheSize 10GB
```


feature » deep into storage space



The main Storage Spaces control panel before anything is created.



A storage pool isn't much use without a few drives to play with.

for a larger number of drives.

In the old parlance, this would have been called a RAID 5 system – striped parity – and it can deal with a single drive failing in the array. A dual-parity option provides parity redundancy, which means two drives can fail with no loss of data; that would have been called RAID 6. Why would you need parity redundancy? Technically, the rebuild time of an array using drives larger than 1TB is so long that another drive could fail, so RAID 6 was recommended in these cases.

A final type of array is just referred to as “simple.” This offers no redundancy – it’s simply pooling drives for storage.

A feature of pooling hard drives that shouldn't be overlooked is the performance through parallel access. Typically, this only improves reads, but it can enhance writes, because the drive controller can queue up parallel reads across the array for both mirrored and parity spaces without additional cache.

The main takeaway here is that drives can be added and removed relatively safely without losing your data, though there is always a “rebuild” time that impacts performance while the parity data is recreated. Windows provides a “optimize” option that should be run when a Space is expanded or reduced. We'll delve into handling Storage Spaces now.

GIVE ME SPACE!

Now we know the basics of what goes on behind the scenes with Storage Spaces, let's move on to how you can set one up within Windows 10 Pro. As we've hinted at, a Storage Space is a “virtual drive,” made up of many other

drives grouped in a storage pool. This means that Windows can build a Space from different types of drives running over ATA, SATA, SAS, USB, or NVMe. It just requires the minimum number of drives for the type of Space you request. You don't even need the full physical capacity available that you request, although Windows throws up warnings as the Space fills up.

For our example, we're going to use four 1TB HDDs in various configurations, and we'll attempt to attach an SSD as a write-back cache via the PowerShell. To kick off, you need to hook up your drives. Once back in Windows, check the drives have been recognized and are blank – ours had a mix of GPT, Dynamic, and previous mirrored arrays on them.

In the search tool, type “diskmgmt” and hit Enter to open the Disk Management console. You have to use your own common sense here to identify your new drives – these need to be cleared before they can be added to a storage pool. Right-click any partitions and select “Delete Volume” – check the contents with “Open” before destroying your data.

With cleared drives, we can move to creating a pool. In search, type “Storage” and select “Manage Storage Spaces.” As none exist, click the link to “Create a new pool and Storage Space.” It presents you with a list of suitable drives that can be added to a storage pool, which should be your blank drives. Select the drives you want to add, and click the “Create pool” button. If you get an error about not being able to complete the operation, try rebooting Windows – we had this issue a couple of times.

Here's where you select the options

for your new Space: name, drive letter, and filesystem. For filesystem, you should only have NTFS available; Microsoft is developing REFS (REsiliant File System), but it's aimed at data center/enterprise applications with checksumming features over mirrored arrays and rapid block re-allocation for virtualisation. It appears Microsoft isn't rolling this out to consumer Windows users – yet.


“Resiliency” refers to the simple, mirror, and parity types we covered earlier. Depending on the type you choose, it alters the overall storage capacities below. As it's a virtual storage system, you can over-allocate the physical space, the expectation being that you'll add more storage when it's required. Once you're happy, click “Create storage space.”

At this point, we're done. Yay! Stop reading, go home. The new Space appears under “This PC” – or whatever Microsoft is calling your PC this year—as a new monolithic drive. You can benchmark, install to, copy from it, and so on, like any normal drive. But it's not a normal drive. Re-open the Storage Space, and you'll see a host of new options, and the status of the overall pool and individual drives.


Because this is all virtualised, you can create additional Storage Spaces within the same pool. Even if the total space exceeds that physically available, Windows doesn't care – it's like that these days. To make any changes, you need to click the Administrator “Change settings” button. As you should see, you're able to create additional Spaces, which takes you through the previous procedure.

You're able to change an existing Storage Space in terms of its name,

“As it's a virtual storage system, you can over-allocate the physical space, the expectation being that you will add more storage when it's required.”



Specify a Storage Space's resiliency and capacity in the control panel.



Once a pool and Space have been created, monitor them here.

drive letter, and overall capacity – its resiliency is fixed. Notice how you don't add drives to a Storage Space; this is restricted to the storage pool. So, if you were expanding your capacity, you'd add a drive to the pool, and increase the available capacity to any Storage Spaces that deserve it. After adding to the capacity of a Storage Space, you should run the "Optimise drive usage" feature – this reallocates data and parity stripes for efficiency over the drives; though be aware that this reduces performance while it's being run.

The "Physical drives" section lists the status and options for all the attached drives. If a drive should fail, or you're planning to upgrade a drive, or for

whatever reason need to remove a drive from the storage pool, you need to click "Prepare for removal." Make a note of the "Model" and "SN" (serial number), so you remove the correct physical drive. Data is moved off the drive, and when finished, a "Remove drive" button appears. The drive won't appear back in Explorer, because it's still formatted as a storage pool device; you need to wipe it through Drive Manager, as we did to prepare the drives at the beginning.

Finally, if you decide Storage Spaces aren't for you, removing them is no harder than clicking the "Delete" button – just make sure you remove all Storage Spaces before trying to delete the pool. In fact, you can't even see the

"Delete" option if there's a Storage Space in use. To reclaim your drives, remove all the Storage Spaces, then you can remove the pool.

That's the easy way to build a straight Storage Space from a storage pool via the GUI, but it doesn't take full advantage of any SSDs, or enable you to increase the write-back cache from the default 1GB. For this, you need to shift to the PowerShell – see the "Solid States" boxout for exactly how that works. Our benchmarks show that it has a dramatic increase to a simple array, boosts mirror performance, and certainly helps the garbage performance that Storage Space parity arrays offer. ■

PERFORMANCE RESULTS

For our testing, we used a bank of four 1TB Toshiba 7,200rpm SATA hard drives in a number of Storage Space configurations: Simple, Mirrored, and Parity. Alongside this we ran two mismatched 480GB SSD SATA drives. We're only

interested in an overview of performance, so we relied on CrystalDiskMark 6 for our results. We even created a couple of Spaces through the Spaces GUI to see what it would make of them, and if the PowerShell route was actually superior.

CONFIGURATIONS			BENCHMARK RESULTS (MB/S)			
SPACE TYPE	HDD TIER	SSD TIER	SEQ R/W	4K Q8T8	4K Q32T1	4K Q1T1
NONE	1X 1TB HDD	NONE	156.9/156.8	1.39/1.11	1.44/1.12	0.59/1.16
SIMPLE	4X 1TB HDD	NONE	164/170	1.70/1.43	1.71/1.54	0.65/1.57
SIMPLE	4X 1TB HDD	2X SSD	1,115/1026	587/627	244/195	34/98
MIRROR	4X 1TB HDD	NONE	176/164	2.38/1.34	2.38/1.24	0.65/1.37
MIRROR	4X 1TB HDD	2X SSD*	216/284	10.7/73.9	10.5/46.7	1.23/42.9
MIRROR	4X 1TB HDD	2X SSD	565/514	462/314	249/99.9	35.7/72.5
PARITY	4X 1TB HDD	NONE	313/65.1	3.97/0.48	3.98/0.38	0.65/0.46
PARITY	4X 1TB HDD	2X SSD*	307/238	4.07/0.29	4.16/1.77	0.90/33.5

*GUI-created Storage Space



System News

Competition is always good for prices, but Mark Williams is wondering when the top end of the GPU town will see competition once again?

As AMD stands toe-to-toe with Intel in the CPU market, grabbing market share hand over fist, the GPU market has been largely left to Nvidia to dominate with recent estimates by the likes of Jon Peddie Research and some of our local retailers putting Nvidia's discrete GPU market share at anywhere between 80% and 85%. Nvidia is the undisputed king of the GPU hill.

Those stats however were from a few quarters ago. Now, in Q3 2019, while Nvidia is still the undisputed top dog, AMD has recently clawed back some market share, which continues to grow, in part thanks to two factors. Firstly, the pricing of its RX 500 series cards has been very aggressive, placing them in a price bracket no Nvidia GPU comes near to touching, bar the GTX 1650, which doesn't even perform as well as the RX 570 but still costs \$40 more. So, in the low end of the market AMD has reduced its prices competitively.

Secondly, AMD finally launched its new Navi-based RX 5700 and 5700 XT cards, and while the reference blower coolers on them are decidedly average due to high fan noise and poor thermals, the AIB partner versions with better coolers have dropped, offering superior thermals and noise levels and performance levels (on the XT model) that slay the RTX 2060 Super and are within spitting distance of the RTX 2070 Super.

One could say that simply offering something competitive at all in a duopoly market is in itself enough to grab easy market share. After all, AMD still hasn't built a GTX 1080 Ti killer, let alone an RTX 2080 Ti competitor. In fairness, the Radeon VII and RX 5700 XT come close to matching the GTX

1080 Ti, but just not decisively (on average).

As good as the RX 5700 parts are on their own merits, AMD is still undercutting Nvidia's closest products and as a result is wiping the floor in the price to performance charts. To illustrate, the cheapest RX 5700 and 5700 XT's sell (at the time of writing) for \$550 and \$640 respectively. Compared to the RTX 2060, 2060 Super, 2070 and 2070 Super which sell for \$570, \$730, \$850 and \$910, respectively.

When you realise that the RX 5700 on average beats the RTX 2060 by 5% yet costs \$20 less, it's easy to see why AMD is picking up market share quickly. The price-to-performance is most pronounced with the RX 5700 XT, which on average beats the RTX 2060 Super also by 5% but costs a surprising \$90 less. The XT also happens to equal the RTX 2070 for performance which costs an amazing \$210 extra. As you can see, the RX 5700 series, and the XT model especially, effectively kills Nvidia's mid-range cards on price and performance.

Nvidia do of course offer unique features (like raytracing) and typically better drivers, which one must consider as well, but these price differences are truly hard to overlook.

There is word going around that "Big Navi" is on the way from AMD, which ratchets up the Navi core counts to maximum with the intent to topple Nvidia's best. Unfortunately, with a long wait before those arrive, AMD must keep playing the pricing game in the mid-range, leaving the high end of the market solely to Nvidia, allowing it to charge whatever it likes.

SHOP TALK

How did AMD graphics card sales perform for you in Q2 this year? With the release of the 5700 series, does it look like AMD will continue gaining market share?

Jaimie, Leader Computers:

"At Leader we carry four brands of AMD cards: Asus, Gigabyte, MSI and the exclusively AMD partner Sapphire. We have seen AMD market share increase across the board, and Sapphire sales increased significantly with the release of the 5700 and 5700XT. Overall the market has performed strongly, with AMD's value proposition hard to deny, just like their CPUs! Gamers have never had so many options for their budget."

John, TI Computers:

"While not quite competing with Nvidia, AMD didn't do too badly in Q2 thanks to the sharp pricing. However, the lack of high-performance GPUs really limits their growth big time. We are pleased to see new GPUs such as RX5700XT and RX5700 offering a similar performance and price point left by GTX1080Ti, which consumers missed the most. Take the Asus Strix RX5700XT as an example, the selling price of \$899 actually made it a no-brainer. I can foresee AMD eating up a big chunk of the market share in the remainder of this year if the strong demand can be fulfilled."

Market Snapshot

A sampling of PC systems available this month.



Mwave S66a

\$3,799 TINYURL.COM/APC473MWV

An interesting build this one. It goes quite high-end on the non-processor parts like motherboard, RAM, cooler and SSD, but then dials it back on the parts that make the most difference to performance: the CPU and GPU. Don't get the wrong idea, it's still a very potent combo for 1440p gaming, but for the price there will be faster options out there depending on your needs. If you're thinking long term, this is a great system for those wanting a solid base to upgrade from later when there are better options available. If you're wanting the best gaming performance straight out of the box at this price, look elsewhere, but lots of RAM and a PCIe 4.0 SSD makes this immediately best for those processing or moving about lots of data as well as gaming.

KEY SPECS: CPU: AMD Ryzen 7 3700X **COOLER:** Corsair H100i **MOTHERBOARD:** Asus X570-E **GRAPHICS:** Asus Radeon RX 5700 XT 8GB **MEMORY:** 32GB Corsair DDR4 Dominator Platinum 3600MHz **STORAGE:** Corsair 1TB NVMe Gen4 M.2 SSD, WD Black 2TB 7200RPM HDD **POWER SUPPLY:** Check with the retailer before purchasing **CASE:** Corsair Carbide 678C.



Scorptec Nemesis RTX 2080 Ti

\$3,699 TINYURL.COM/APC473SCP

For all the hype surrounding Ryzen, when it comes to gaming AMD still can't topple two of Intel's best: the 9900K and 9700K. This suitably titled Nemesis build serves up some of the best performance you can find for gaming that AMD currently simply can't match. The 9700K will aid in serving up copious amounts of frames for those searching for 250Hz+ gaming and paired with the current best GPU in the RTX 2080Ti, this system will easily handle it. Water cooling will aid with overclocking the CPU to further increase FPS in games if it strikes your fancy. The RAM and SSD are lower spec than what else is out there, but if your primary usage is gaming, these are perfectly fine.

KEY SPECS: CPU: Intel Core i7 9700K **COOLER:** Cooler Master MasterLiquid ML240R RGB **MOTHERBOARD:** MSI MPG Z390 Gaming Edge AC Motherboard **GRAPHICS:** MSI Ventus GeForce RTX 2080 Ti OC 11GB **MEMORY:** Corsair Vengeance LPX 16GB 3000MHz DDR4 **STORAGE:** Intel 660P 512GB M.2 SSD, Seagate BarraCuda 2TB HDD **POWER SUPPLY:** SilverStone Essential 650W **CASE:** Fractal Design Meshify.



JW Computers Intel NUC BOXNUC8I3BEH4

\$449 TINYURL.COM/APC473JWC

This system, comparable to the Asus MiniPC from TI computers below, is quite a bit more powerful on the processor front, with twice the thread count (4), a 50% higher base clock (3.0GHz) and an almost 40% higher turbo clock (3.6GHz) with Iris Plus integrated graphics included, it is a nice boost in performance. There is a big catch though, this is a barebones system – meaning BYO memory and mass storage. Then there's a Windows license on top if you're wanting that, too. All of which will substantially add to the total price of the system. Great for those that'll be doing work rather than just media consumption and need the responsiveness in a tiny form factor.

KEY SPECS: CPU: Intel Core i3-8109U **COOLER:** OEM **MOTHERBOARD:** OEM **GRAPHICS:** Integrated **MEMORY:** None (supports DDR4-2400 SO-DIMM) **STORAGE:** None (supports M.2 and SATA) **POWER SUPPLY:** Included power brick **CASE:** OEM.

TI Computers Asus MiniPC PN40

\$499 TINYURL.COM/APC473TIC

Have you the need for a very basic computer? Be it for just data entry or simple media playback or just web browsing? This system takes on the NUC class of products with a very basic list of specifications. The dual core non-hyperthreaded CPU at its heart is the biggest constraint for performance but what it lacks in grunt it makes up for with efficiency, sipping just 10W allows this system to run completely fanless for silent operations. Amazingly, it ships with Windows 10 Pro, which alone can retail for \$200+. With good video output support for its size (HDMI and Mini-DP at up to 4K) at such a low cost it's hard to beat if all you need is to show/view something on a monitor or two.

KEY SPECS: CPU: Intel Celeron J4005 **COOLER:** OEM **MOTHERBOARD:** OEM **GRAPHICS:** Integrated **MEMORY:** 8GB Kingston DDR4 2666MHz **STORAGE:** 250GB WD SSD **POWER SUPPLY:** 65W OEM **CASE:** OEM.



PC blueprints

The APC team's picks for a part-by-part perfect PC build to suit your budget.



PARTS LIST

PART		PRICE
CASE	BITFENIX NOVA TG	\$49
PSU	CORSAIR VS450 80+ WHITE 450W NEW	\$49
M/BOARD	ASROCK B450M R4.0 AM4	\$95
CPU	AMD RYZEN 5 2600X	\$229
GPU	ZOTAC GAMING GEFORCE GTX 1660 6GB NEW	\$359
RAM	8GB (2X 4GB) TEAM T-FORCE VULCAN Z @ 3,000MHZ NEW	\$79
SSD	120GB CRUCIAL BX500 2.5-INCH	\$29
HDD	1TB WD BLUE 3.5-INCH HDD	\$59
OS	UBUNTU DESKTOP LINUX 18.04 LTS 64-BIT	\$0

APPROXIMATE PRICE: \$948

We played with the idea of putting some new toys into the budget build this month, what with the advent of the Ryzen 3000 series of CPUs, but ultimately we opted to keep the cost down. New X570 motherboards, which are really needed to get the most out of those processors, are expensive; while the trusty Asrock B450M R4.0 just keeps getting cheaper. We also swapped out the GTX 1660 GPU from Gigabyte to Zotac to preserve the price, and last month's Kingston HyperX RAM is no longer the best value, so back to Team Group's T-Force line we went, sticking with a two-piece kit that would offer comparable performance. Lastly, we swapped the EVGA power supply for the remarkably priced Corsair unit listed. At just \$49 it's exceptional value and has all the power this rig needs.



PARTS LIST

PART		PRICE
CASE	FRACTAL DESIGN MESHIFY C	\$149
PSU	500W COOLER MASTER MWE 500 SILENCIO 80+ BRONZE NEW	\$55
M/BOARD	ASROCK X570 PHANTOM GAMING 4 NEW	\$259
CPU	AMD RYZEN 5 3600X NEW	\$385
COOLER	AMD WRAITH SPIRE NEW	\$N/A
GPU	SAPPHIRE RADEON RX 5700 8GB NEW	\$699
RAM	16GB (2 X 8GB) CORSAIR VENGEANCE LPX @ 3,200MHZ	\$139
SSD	250GB SAMSUNG 970 EVO PLUS M.2 PCIE	\$119
HDD	1TB WD BLUE 3.5-INCH HDD	\$359
OS	WINDOWS 10 HOME 64-BIT OEM	\$140

APPROXIMATE PRICE: \$2,304

Big changes ahead this month for our mid-range build. The first step was to upgrade the processor, going straight for the high-quality Ryzen 5 3600X, with a hefty base clock of 3.8GHz. These new 3000-series CPUs aren't too power-hungry either, meaning we could safely dial back the PSU to Cooler Master's 500W Silencio. Although this won't leave too much additional power for overclocking, the Ryzen 3000 processors don't have much overclocking headroom anyway. To support the new CPU, we needed a new motherboard: the X570 Phantom Gaming 4 from Asrock, one of the best value X570 boards available right now. We're opting for the stock Wraith Spire cooler here to keep the price down, but if you've got the money to spare, an AIO cooler isn't a bad idea. Lastly, we've replaced the previous RTX 2060 with AMD's RX 5700 GPU.



UPGRADE OF THE MONTH

**EVGA GeForce RTX 2070 Super
Black Gaming 8GB**
\$889, WWW.EVGA.COM

The new Super cards are, to put it bluntly, just good products. They're not particularly exciting, but what they set out to do, they achieve with gusto. The improvements over their original counterparts are modest, but significant enough to make a difference, and this was achieved without raising the price tag, too. The 2070 Super is far closer in both performance and internal design to the original 2080 than the 2070, and that really makes a difference in game.

"The new Super cards are, to put it bluntly, just good products."

A wide selection of changes this month. Most crucially was a new GPU: the RTX 2070 Super. Thanks to its upgraded architecture, this new card provides performance only marginally beneath the original 2080 for a significantly lower price. We've gone with a twin-fan offering from EVGA, costing a measly \$889 for a 1440p powerhouse. The RTX 2070 Super is so good that we even considered putting its little sibling the 2060 Super in our mid-range build, but if we were planning on raising the price by that much, we'd have chosen the Radeon RX 5700 XT instead.

There wasn't much need to change the processor; while we considered the new Ryzen 9 3900X, the fact is that Threadrippers are simply more powerful. As such, the motherboard didn't need re-evaluating either, although we did spring for a better AIO cooler from Fractal Design for the 2950X (albeit a less flashy one).

Lastly, the dawn of Gen4 M.2 drives has led to a pleasing drop in prices for previous-generation NVMe SSDs. This may not last, so if you're looking to upgrade your storage, now might be the best time to do so. We decided to upgrade to the 970 Evo Plus, netting us superior SSD performance without having to raise our storage budget. At \$160 right now, the 970 Evo Plus is an absolute steal. ■

PARTS LIST

PART		PRICE
CASE	NZXT H700I	\$219
PSU	850W CORSAIR TX-M SERIES 80+ GOLD	\$179
MOBO	ASUS PRIME X399-A	\$419
CPU	AMD THREADRIPPER 2950X	\$1,179
COOLER	FRACTAL DESIGN CELSIUS S36 BLACKOUT 360MM AIO	NEW \$296
GPU	EVGA GEFORCE RTX 2070 SUPER BLACK GAMING 8GB	NEW \$889
RAM	32GB (2X 16GB) G.SKILL RIPJAWS @ 3,200MHZ	\$299
SSD	500GB SAMSUNG 970 EVO PLUS M.2 NVME	NEW \$160
HDD	2X 3TB SEAGATE BARRACUDA COMPUTE	\$238
OS	WINDOWS 10 HOME 64-BIT OEM	\$140
APPROXIMATE PRICE:		\$4,018



Video-Editing Powerhouse

In need of a new rig for editing video? Look no further...

Inspired by a handful of reader letters, this issue we're building a video-editing machine. The most dedicated readers among you might recognise the case as the same one we used for a previous video-editing build, but this time the internals are very different. With AMD releasing a bevy of new components, we knew it was time to put some of them to the test in a new rig. We also tried to stick to a relatively sensible budget for a workstation PC, setting ourselves a cap of \$3,500.

Without blowing closer to \$4,000, we won't be making a godly beast of a computer dedicated to 4K video editing. What we're producing is an effective but not overpriced workstation, ready for all manner of office work, as well as video editing, image manipulation, and audio work. Flexibility was key: A system that could carry out complex

tasks quickly and effectively. There was a dual purpose, though: Testing the capabilities of the array of new hardware from AMD. While we've played around with individual components, this is our first build properly featuring the new hardware; that means a next-generation CPU and GPU from AMD, as well as a new motherboard and drive, to produce a build capable of taking full advantage of the new NMVe Gen4 and PCIe 4.0.

For this rig, as with all PCs built with editing in mind, we needed a powerful CPU, capable of performing processor-intensive tasks such as rendering and video encoding. For that, we're using one of the more powerful new 3000-series Ryzen CPUs: the Ryzen 7 3700X, with a base clock of 3.6GHz, and a potential boost clock of 4.4GHz. Let's see how this brand new system turns out.

READY TO EDIT

While the Phanteks Eclipse P600S can come with some flashy LED lighting options, we're forgoing all that and sticking with the simple, clean-cut appearance of the white model. That's not to say we're abandoning all lighting whatsoever, mind you; the Ryzen 7 3700X's stock cooler comes with some flashy RGB effects, and we're pairing it with some Corsair Dominator Platinum RGB, a four-piece kit that comes with 48 individually addressable LEDs. Yikes. So, lighting is pretty much sorted. That's good, given that the Gigabyte X570 Aorus Master only has some very subtle orange lighting on one edge, and the Radeon RX 5700 has no lighting whatsoever.

We like the look of the RX 5700, in fact. Unlike the odd dented appearance of the 5700 XT, the standard 5700 model from AMD is a sleek gray oblong of brushed metal that reminds us of older cards. It's a quality product that doesn't cost too much, and while users handling extremely high-res footage (think 8K) will want something more powerful, it'll comfortably handle most editing work.

Adding to that is a slick new NVMe Gen4 M.2 SSD, Corsair's 1TB MP600. The new X570 motherboards are prepared to handle massively faster M.2 transfer speeds with NVMe Gen4, so for a system that will likely be transferring large files on a regular basis, this is ideal. Surprisingly enough, this system won't draw too much power without some pretty savage attempts at overclocking, so we're sticking in one of Fractal Design's brand new PSUs. The Swedish manufacturer has danced with power supplies before, providing its Integra PSUs (integrated power units, in case the name didn't give that away) for several of its cases. Its new Ion+ series are high-quality modular PSUs, available in a variety of wattages. We're using the 660W model, because it provides enough power for our system without breaking the bank.




1 DO OR DISASSEMBLE

This might be the easiest case to prep we've ever come across. The side panels – one glass, one metal – open with magnetic clasps, and swing out on hinges, from where they can be easily lifted off. Meanwhile, the front and top panels are secured solely with magnets. Make sure the cables from the front I/O and the PWM fan hub (at the top) are nice and tidy in advance, using the Velcro straps Phanteks has provided. We're using all the stock case fans, so there's no fiddling around to be done. Feel free to remove any drive cages or mounts for HDDs and SATA drives, as we're sticking to M.2 storage for this build. Lastly, make sure your build area is clear and you have all the tools you need: a Phillips head screwdriver, thermal paste, and a bowl for your screws.

PARTS LIST

PART		STREET PRICE
CASE	PHANTEKS ECLIPSE P600S WHITE	\$220
MOTHERBOARD	GIGABYTE X570 AORUS MASTER	\$629
CPU	AMD RYZEN 7 3700X	\$519
CPU COOLER	AMD WRAITH PRISM (STOCK COOLER)	N/A
MEMORY	CORSAIR DOMINATOR PLATINUM RGB 4X 8GB @ 3,600MHZ	\$595
GPU	AMD RADEON RX 5700	\$529
PSU	660W FRACTAL DESIGN ION+	\$189
STORAGE	CORSAIR MP600 FORCE SERIES GEN4 1TB M.2 SSD	\$445
OS	WINDOWS 10 HOME 64-BIT	\$140
TOTAL		\$3,266



2 THE YEAR 3000

The P600S is quite bulky, so we'll prep the motherboard assembly outside the case. Like many of the new X570 boards, heat shielding and dedicated cooling mean that the Aorus Master is pretty heavy, so be careful when removing it from the antistatic cover. Unless you've got a dedicated workspace, use the box it came in as a surface to mount it as you add components. Start with the processor: Simply lift the retention arm, line up the gold corner of the chip with the triangle on the bracket, drop it into place, then lower the arm so it locks in position. Moving on, this mobo has four DIMM slots; that's exactly how many sticks of RAM we're using – how convenient. Release the clips at either end of the slots, then push the sticks into place until they click on both sides. All done!




3 HIGH-SPEED PURSUIT

How we love the NVMe Gen4 M.2 drives. The transfer speeds are monstrous. What we're not so sure about is manufacturers slapping (often ugly) integrated heatsinks on them, when the X570 boards that support them almost unilaterally have their own heat management solutions. Such is the case here, although the Corsair MP600 looks very inoffensive. Unscrew the smallest of the Aorus Master's three M.2 heat spreaders and remove it, then fit the MP600 drive into place, and screw it down. Put that loose heat spreader somewhere you won't lose it! If you'd prefer to use it, the MP600's heatsink can be removed by pushing in the two clasps on either side. Take off the plastic cover on the sticky pad beneath the heat spreader you removed from the mobo, and ensure the M.2 drive is properly attached before you screw it back in.




5 RADEON-ATION

Be warned: If you have gargantuan paws like our six-foot-three staff writer, you may wish to skip to step six before installing the GPU. While we usually try to get all the components fitted before messing with cable management, we found that squeezing in some of the cables from the front I/O proved tricky once the GPU was in place. The Radeon RX 5700 is lovely looking, though; it feels solid and fits nicely in the case. Remove two of the blanking plates at the rear of the case. Make sure the PCIe slot's clasp is open, then push the card into it, so it clicks into place, before securing it to the backplate with the screws that had secured the blanking plates. The RX 5700 is heavy, so secure it tightly, and be careful not to drop it on your board, lest you damage other components.



4 PRISMATIC

With most of the components in place, the next step is to install the motherboard. This is pretty straightforward; lay the case on its side, then lower the board in before screwing it down, being careful to fit the rear I/O beneath the fitted case fan. As is often the case, you may have some difficulty, so if the fan gets in the way, remove it and reinstall it to the rear of the case once the motherboard is secured. Now get the Wraith Prism cooler ready. Your board should come with two AMD stock standard clips fitted either side of the CPU socket – if not, fit those first. Once affixed, apply a small blob of thermal paste to the CPU (we're using Noctua NT-H1 compound) before lowering the cooler into place, and making sure both clips are hooked to the bracket. Pull the clasp so it locks into place, and your cooler is secured.



6 MISCHIEF MANAGED

Lastly, we install the PSU. The Ion+ is modular, so you need to sort out which cables you require. You need the thick 28-pin connector for the mobo, and two eight-pin connectors for the CPU. The GPU takes one eight-pin and one six-pin connector; you can find both on one cable, in the form of two six-pin and two two-pin connectors. For the eight-pin GPU slot, fit one of the two-pin connectors to the side of the attached six-pin one. Lastly, you need one SATA accessory connector, for powering the PWM fan hub. Next, fit the PSU at the rear and run the cables through to the sockets on the board. The P600S has plenty of slots and Velcro straps for cable management, so tidy up the cabling to the rear of the motherboard plate before sealing this big white box shut.



1 While it might only be the stock cooler for the 3700X, the Wraith Prism is actually very effective when it comes to keeping your processor cool. It looks pretty good, too, and doesn't impinge on the DIMM slots, even with bulkier RAM.

2 The RX 5700 is a great mid-range GPU, and there's always room for upgrading in this build. We're still waiting on a high-end card from AMD, though....

3 The PSU shroud has plenty of space to conceal any loose cables, but thanks to the modular design of the Ion+ 560P, there shouldn't be too much of that nonsense to deal with.

A quality rig with room for improvement

We're pretty pleased with the outcome of this build. It might not be as optimised for gaming as some of our "Blueprint" PCs – though it should still handle most games at 1440p – but it smashed our benchmark zero-points in other areas. The Gen4 M.2 drive is pulling its weight and then some, getting close to a dizzying 5GB/s of read speeds and over 4GB/s of writes. Meanwhile, the Ryzen 7 3700X and Radeon RX 5700 both did great, performing well in Cinebench and 3DMark.

Were price no object, or if we were focusing on extremely high-end performance, we'd have spruced up

some parts of the build. This system could comfortably support a high-end GPU, such as a Radeon VII or an RTX 2080 Ti, without starting to struggle, although for extreme 8K video editing, it may be worth upgrading to 64GB of RAM, too. As a mid-level workstation, though, this rig is easily capable of doing the work. The 3700X is a quality CPU; we did contemplate using a Ryzen 9 3900X, but the price increase against performance didn't seem worth it.

The Gigabyte X570 motherboard is a great, dependable board for supporting a flexible build, and likely won't need upgrading for a long time. The Eclipse

P600S also leaves plenty of room for expansion, with space for a larger GPU or a soundcard. There's also room for an all-in-one cooler; a worthy extra if you're looking to overclock. The Wraith Prism is actually a fantastic cooler, considering it comes with the 3700X, and it can comfortably handle the CPU's automatic overclocking functions without overheating. Ryzen 3000 processors have some issues with manual overclocking, but if you've got the expertise to make it work, a more robust cooling solution may be required.

These builds often serve as more of a guideline than a hardline tutorial, so we expect you to swap out components here and there. One part we recommend, though, is the Fractal Design Ion+ PSU. Its modular design and clean appearance make it a very attractive choice of power supply. It also comes in 760W, and 860W models, should your system need a little extra juice.

It's hard not to like this build. The setup was straightforward and the end result is a pleasingly simplistic system that runs well and looks good; the RGB lighting is restrained, but flashy enough to impress. The only other thing we'd consider adding is a hard drive, because our MP600 only packs 1TB of storage. It might sound a lot, but for storing large video files, slapping in a roomy HDD might become a necessity. ■

BENCHMARK RESULTS

	ZERO-POINT	THIS SYSTEM
CINEBENCH R15 MULTI (INDEX)	959	2,096 (119%)
CRYSTALDISK QD32 SEQUENTIAL READ (MB/S)	3,442	4,998 (45%)
CRYSTALDISK QD32 SEQUENTIAL WRITE (MB/S)	1,706	4,286 (151%)
RISE OF THE TOMB RAIDER (FPS)	95	84 (-12%)
TOTAL WAR: WARHAMMER II (FPS)	55	55 (0%)
TOM CLANCY'S GHOST RECON WILDLANDS (FPS)	48	51 (6%)
3DMARK: FIRE STRIKE (INDEX)	9,128	20,755 (127%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of an Intel Core i5-8400, 16GB (2x 8GB) of Corsair Vengeance LPX @ 2,666MHz, an MSI GeForce GTX 1080 Armor, and a 500GB Samsung 960 Evo M.2 PCIe SSD. All tests were performed at 1440p at the highest graphical profile.

how to

» QUICK TIPS

Experts solve your computing problems

APC and its readers can be one giant helpdesk. If you have a technical problem, chances are one of us can solve it.

APPLICATIONS

MISSING INSTALLER BLOCKS UNINSTALL

I'm unable to remove a program from my PC because I'm told the installer is missing. How do I go about resolving this issue?
Jack Mahoney

Start by visiting the program's home page – download the latest installer from there and then run it to see whether it detects the existing installation and allows you to remove it. If not, and the version number is roughly the same, you may find installing the program over the top of itself restores the missing links and allows you to subsequently remove it. If this fails, then download and install IObit Uninstaller 8.5 Free (www.iobit.com/en/advanceduninstaller.php). Launch the program and click the toolbox icon – if the unwanted program is currently running, choose the Easy Uninstall button, then click and drag

the crosshairs button over the program window, selecting Uninstall from the pop-up menu that appears. If this fails, click the toolbox and choose 'Force Uninstall' instead, clicking Browse to select the program's main executable file from either the Program Files or Program Files (x86) folders. From here IObit Uninstaller will attempt to remove the program manually and clean up as many leftovers as possible.

Nick Peers

MAC


MICROSOFT OFFICE 2016 IS SLOW TO START

After upgrading to Microsoft Office 2016 (from 2011), opening each app has been painfully slow, with icons bouncing in my MacBook's Dock for 20 seconds or more. Once running they're fine. How can I fix this?


Deane Robinson

The apps in Microsoft Office 2016 do tend to be more sluggish to start up, but

"Start by visiting the program's home page – download the latest installer from there and then run it to see whether it detects the existing installation and allows you to remove it."



IObit Uninstaller may help ditch apps that 'lost' their installer.



When your fonts have gone wonky and nothing else can fix them, you can use Font Book to restore the standard set.

shouldn't be taking that long. Word can sometimes struggle to find old folders listed in the File Locations section of its preferences, but delays like this most commonly result from font problems.

It's usually best to put the special Microsoft fonts in a folder named Microsoft in /Library/Fonts, although Office 2016 now stores a copy of its bundled fonts internally too. Check that these are available in Font Book. You may need to clear font caches. The simplest method is to start up in safe mode (see bit.ly/mfsafemd). Let your Mac get going fully, then restart back in its normal mode. You may also need to open Font Book again and choose File > Restore Standard Fonts to return your Mac's fonts back to the set that ship with macOS. This also moves any extra ones you've installed to a separate folder, disabling them. You'll then need to move the Microsoft Office fonts back from /Library/Fonts (Removed), which is created during that restoration.

If all else fails, uninstall Microsoft Office, then reinstall it using a freshly downloaded copy. If you're still mystified as to what's causing this, try EtreCheck's generic diagnosis tool (eclecticlight.co), which can spot other clashes and issues that might be responsible for generally poor performance.

Howard Oakley



Orico tolless enclosures are cheap, but they're well-built and easy to install


RECOVERY ACCESS DATA FROM DEAD PC'S HARD DRIVE

I feel a bit of a fool. My ageing laptop has suddenly stopped booting and I need to recover data from the hard drive having not recently backed up. Is this possible, and if so how?

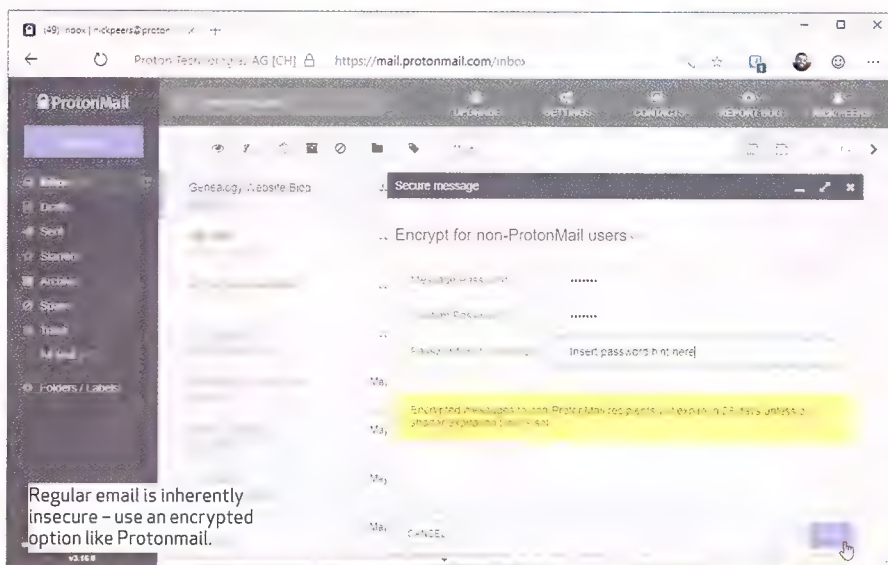
Thomas Lewisham

Thomas revealed the laptop wouldn't even boot as far as the initial start-up screen, which gave us hope the drive inside was both recoverable and reusable. The laptop was Thomas's secondary machine, and he confirmed he'd like to continue using the 500GB

drive as backup, so we pointed him in the direction of an inexpensive 2.5-inch drive enclosure (search eBay for 'Orico 2.5-inch enclosure' - currently around \$12). He then physically removed the drive from the laptop, mounted it in the new enclosure and connected it to his desktop PC. The drive appeared in Windows, complete with all data intact. Thomas encountered an 'Access Denied' error attempting to open his user folder - clicking 'Continue' resolved this problem by giving his user account both read and write access to the drive. We also advised him to take a fail-safe drive image backup of the laptop drive



Store all your Office macros in a dedicated document for easy access.



Regular email is inherently insecure – use an encrypted option like Protonmail.

"Email messages aren't encrypted during transit, and are passed between multiple servers as they journey from sender to recipient. Beefing up security requires action on both your part and your recipients, however."

using Macrium Reflect, which he stored on his network hard drive. Once done, he was free to format the drive for File History use.

Matt Hanson

WINDOWS

CAN'T UPDATE WIN 10

For reasons I'm not sure, any attempt to install the 1903 update results in a "This PC can't be upgraded to Windows 10" error.

Alan Hart

This rather unhelpful error suggests

your PC is permanently incapable of running the Spring Update, but in actual fact it's down to the fact that Windows has detected there's a USB drive or SD card plugged in. Long story short, the 1903 Update has been known to reassign drive letters after the update has been installed, which might cause problems elsewhere. As a result, Microsoft has put a so-called block on the update. The solution is simple: unplug all your external drives, then attempt to install the update using <https://support.microsoft.com/en-gb/>

help/3159635/windows-10-update-assistant. Once it's successfully installed and you're returned to the Windows desktop, reconnect your drives one at a time. We recommend plugging them in to match their current drive letter order, so start with drive E, then move on to drive F, and so on and so forth. This should ensure they're allocated the same letters.

Catherine Ellis

OFFICE

STORE MACROS OUTSIDE OF DOCUMENTS

I'm a big fan of Office's time-saving macros – particularly those I can copy and paste as VBA code from reputable online sources. For security reasons, I'm happy that Word, PowerPoint and so on don't save them into regular documents, but it can be annoying having to recreate them every time I need them. Can you suggest a more elegant way of keeping these macros on tap?

Nigel Mitchell

Word, Excel and PowerPoint allow you to apply macros in one open document to any other. The solution, therefore, is to take advantage of Office's support for macro-enabled documents (.docm, .xslm and .pptm files respectively) – create a new blank document, switch to the ribbon's Developer tab to open Visual Basic, then choose Insert > Module to create a blank module into which you copy and paste all the macro code you need. Close the Visual Basic window, then save it as the relevant macro-enabled document.

Simply open this document whenever you want access to the macros inside – you'll be able to run them in other documents via the Developer tab's Macros button.

Ian Sleightholm

APPLICATION

HOST FAMILY HISTORY WEBSITE ONLINE

I'm looking for a simple way to securely share my family tree on the internet without putting it in the hands of a subscription service or one that may index it for others to find. Can you recommend anything? I'm happy to pay a one-off fee.
Mark Chadwick

I suggest purchasing a copy of Heredis 2019 (around \$60) – not only is it a fully formed genealogy app in its own right but it has an option for publishing your tree via its own free online service – just import your research for another program via the GEDCOM format. Once published, your tree can be opened up for public access or you can password-protect it so only invited guests are able to connect. What we particularly like about Heredis is the slick, easy-to-navigate interface provided for your site visitors, while uploading the latest

changes from your research is also simple too. You can road-test the program before purchase – a fully functional free version with support for up to 50 individuals can be downloaded from the main Heredis website (www.heredis.com/en).
Nick Peers

SECURITY

SEND EMAIL SECURELY

I've heard email is inherently insecure. What steps can I take to remedy this?
Claire Harradine

It's true: email messages aren't encrypted during transit, and are passed between multiple servers as they journey from sender to recipient. Beefing up security requires action on both your part and your recipients, however. You can add encryption plugins for existing mail providers and applications – for example, Thunderbird users must install GnuPG (<https://gnupg.org/>) followed by the Enigmail add-on, while Gmail users can examine the FlowCrypt (<https://flowcrypt.com/>) extension. These options don't encrypt your mailbox, however – to fix that requires switching to a new email service. To examples are Protonmail (<https://protonmail.com/>) and Tutanota (<https://tutanota.com/>). End-to-end encryption applies when conversing with other users of the same service, while emails sent to other addresses are protected by sending an invitation message in their place: recipients must click the link and enter a password (share this separately) to decrypt and read the message in their web browser. Protonmail also supports OpenPGP for those who want to use that to communicate directly with fellow OpenPGP users.
Catherine Ellis

OFFICE



MERGE COLUMNS AND ROWS IN WORD TABLES

I've created a table where several cells have been incorrectly split into columns. Is there a way to merge these without affecting the rest of the table?
Len Steel

While it's possible to manually merge multiple cells easily enough (select them, right-click, and choose 'Merge selected cells') there's no one-click option for rows or columns. Instead, you'll need Word's Eraser tool.

Click inside your table to make the Layout tab (Design tab in Word 2007 and 2010) visible on the ribbon. Select this and then click the Eraser button in the Draw group (Draw borders in Word 2007/2010).

Click and hold the mouse, then carefully drag down the line of the



column or across the row that you wish to erase between the cells you wish to combine. You should see them turn to red as each one is selected.

Sometimes you'll select the wrong elements – try moving the mouse back to undo the selection or click to perform the merge before pressing Ctrl + Z to undo. You can delete individual lines too – just click each one with the eraser.

Once the cells have been combined, their contents will be placed on separate lines within the cell. You can then tidy these up as you see fit after clicking the Eraser button on the ribbon again to de-select the tool.
Carrie Marshall

EMAIL

WHEN TRASH BECOMES IMPORTANT

My Important mailbox has filled up with all the messages that I have been moving to the Trash mailbox, whether for my iCloud or Gmail account. What's causing this?
Ian Hurdley

'Important' isn't so much a mailbox as a label that's being incorrectly applied to your messages when you delete them. You need to discover which setting is labelling your deleted messages as being important, so that they can go into the Trash properly.

First, check your Gmail account's setup in the service's web-based interface. If you use the account on an iOS device, go to Settings > Passwords & Accounts > [Your Gmail account] > Account > Advanced. Under the heading Move Discarded Messages Into, ensure the Deleted mailbox is selected. Also check Mail's preferences on your Mac. Click the Accounts tab, select the account on the left, then click the Mailbox Behaviours tab. You'll see a Trash Mailbox setting; set the adjacent pop-up menu to the corresponding account's Deleted Messages mailbox. Check you don't have anything set that moves or labels messages in the Rules tab in Mail for Mac, in Mail's preferences at iCloud.com, or a filter in Gmail's settings (see bit.ly/gmfilters).
Howard Oakley ■

Get the most from your AMD CPU

Christian Guyton shows us the quick and easy way to overclock a new Ryzen CPU.

Overclocking is a finicky business: Adjusting frequencies and voltages to try to eke out a bit of extra performance from your rig – while simultaneously avoiding the terror of the blue screen of death. It's often a lengthy process, too, demanding repeated trips into your BIOS to make minuscule adjustments. AMD processors are also known to be a little tougher to effectively overclock than Intel's K-variant chips, which have excellent overclocking headroom.

Thankfully, there's an easier way to manage overclocking on AMD processors. While there is plenty of licensed and unlicensed overclocking software out there, AMD has its own official piece: Ryzen Master, which can adjust both CPU and RAM settings on the fly in Windows. No BIOS trips necessary.

Ryzen Master isn't necessarily the easiest piece of software to use, mind you. The overclocking system is fundamentally the same as overclocking in the motherboard BIOS, insofar as it involves raising the operating frequency of the CPU, then raising the voltage supplied to the motherboard socket in order to prevent crashes. However, the Ryzen Master interface is mercifully much easier to navigate than most BIOSes, and it comes equipped with a few other nifty features that make overclocking a breeze, even for those without the usual tech mastery required.

1 GET PREPPED

First, you need to get set up. Before doing anything, take a trip to Google and find out the actual overclocking headroom on your CPU. Overclocking is best done gradually, but it's good to know what the effective maximum clock you're likely to reach is, so you don't push too far too early, and risk damaging your processor. It's also worth checking purely because there's a chance that your CPU isn't actually up to much when it comes to overclocking – lower-end chips can have very minimal potential for it, and sometimes it's best to just stick to automatic overlocks. We'll get to that later, though.

You're also going to want a fairly robust cooling solution. That means no



stock air coolers, although AMD's Wraith Max may cover you for some tinkering. Ideally, we recommend an all-in-one cooler with a 240mm radiator – or larger, if you're using a custom liquid-cooling loop to cool multiple parts. Once your system is ready, download and install the Ryzen Master software from the AMD website, and fire it up [Image A]. It's time to get started.

2 RUNNING ON AUTOPILOT

Upon booting up Ryzen Master, you're greeted with the homescreen. This displays live data from your CPU, including the current operating frequencies of each individual core, the power draws, and the temperature. The first step here is to choose a profile that you're going to work in. Ryzen Master comes with two preset profiles, called "Creator Mode" and "Game Mode," then two blank profiles for you to name yourself. All four actually have the same initial settings, but for now select "Profile 1."

Check out the annotated diagram to make sure you're on the right page. From here, there's a few different options. If you like the easy life, you can hit the "Auto Overclocking" option at the top of the screen, then select "Apply and Test." Ryzen Master has its own benchmarking test, which it can run after applying new settings to stress-test the CPU. However, we've found that this test isn't too rigorous –


hardcore benchmarks such as Prime95 and AIDA64 put the processor under heavier strain, and in our testing, the CPU would regularly pass the Ryzen Master stress test before crashing within two seconds of Prime95's operation.

With the auto overclocking mode, you have a couple of settings to tweak. Power supply options can be changed at the top of the interface, although the only setting we really recommend changing is "Boost Override CPU." This should default to 100, but if you're not seeing the performance you want after engaging auto overclocking, try turning it up by increments of 25. The other potential option here is the "Precision Boost Overdrive," originally created by AMD for Threadripper processors. This supposedly redistributes power with even greater precision, as required by individual system components and CPU cores, but ultimately doesn't produce much of a difference in performance from the auto overclocking mode.

3 CLOCKED IN

If you're up for a bit more of a challenge, though, there's the manual mode. If you're familiar with CPU overclocking via the BIOS, this stage should be pretty straightforward. Don't worry if you're not, though; we're going to give you the full rundown. Start by selecting manual mode, then make sure that all the cores are linked

A Quick Guide to the Ryzen Master UI



1 CHOOSE YOUR MODE

You can select between modes here. All profiles are set to "default" at first; no overclocking whatsoever.

2 MAKE THE LINK

These icons indicate whether each CCD and CCX (core groups) are linked. Green is linked, red is unlinked. Linked cores match each others' frequency.

3 FIDDLE WITH THE FREQUENCY

This is the operating frequency of each core, adjustable by the green bar. Cores can be disabled by clicking the green box, but we wouldn't recommend this.

4 SPEED STARS

This icon represents which core is the fastest: Yellow stars denote the fastest core; gray is the fastest within that CCX.

5 AMP UP THE VOLTAGE

Voltage control should always be left included. Tweak it upward in small increments using the arrows.

6 FORGET THE MEMORY

Memory control is excluded, and set to auto by default, and unless you know what you're doing, it's best to leave it that way.

by clicking the icon on the right-hand side, so that all the cores are green. (Check our annotation above if you're not sure which icon we're talking about.) From there, you can either click and drag the green slider bars to adjust the operating frequency, or click the number and type it in directly.


Generally, it's best to raise the frequency in increments of 0.1GHz between tests. If you know the maximum boost clock of your processor, though, it's a safe bet to start by stepping up the frequency by half the difference between the base and boost clocks. So, if your processor has a base clock of 3.7GHz and a boost clock of 4.3GHz, you can cut out some of the testing and go straight to 4.0GHz without too much concern of damaging your processor. Once your chosen frequency is set, hit "Apply and Test." This may require a restart – allow the system to reboot, and Ryzen Master will begin automatically after a few moments. As we mentioned earlier, it's best to run a more demanding program than the Ryzen Master stress test, just

to make sure that the system will run. Chances are, you'll experience a crash here; don't worry – that just means that you'll need to raise the voltage being supplied to the CPU. Do this in Ryzen Master, first by 0.2V and then by 0.1V at a time, until the system can withstand a sustained stress test at your overclocked frequency without crashing.

4 STABILISING

Once you've got a stable overclock, keep stepping it up by increments of 0.1GHz. When your system crashes under stress, raise the voltage by a small amount, then continue. Ryzen Master insists on auto-adjusting the voltage to more precise amounts, but don't worry about this. Eventually, you will reach a stage where operating the system under stress becomes impossible, regardless of the voltage increase. At this point, dial the frequency back by 0.1GHz, and find the lowest voltage at which the system can be successfully stress-tested at – that's your stable overclock. All done!

Of course, there are more precise ways to overclock. Your CPU – particularly one of the higher-end 3000-series Ryzen chips – will demonstrate that certain cores operate at better speeds than others. Ryzen Master highlights the faster cores, so if you're struggling to get a functioning overclock, focusing on individual cores can potentially solve your issues. To do this, unlink the cores so that all the icons are red, then clock up the fastest cores, and raise the voltage where needed. The performance difference may not be as noticeable, but these faster cores can potentially be clocked higher than the overall maximum attainable clock speed for the CPU using all cores, which results in superior performance in tasks that only use a single core. ■



Sound advice for iMovie

How to enhance your project's sound using George Cairns' recording and editing tips.

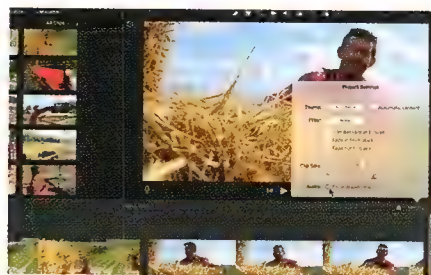
Gadgets such as GoPro cameras and the iPhone can do an excellent job of capturing high-quality video footage for your iMovie projects. However, the audio content you end up with may not match the quality of the visuals. For example, sound levels can be weak because of the distance between the camera's microphone and the subject.

Background noise, such as a passing car or plane, may also drown out your subject's voice.


Luckily, iMovie boasts a collection of audio post-production tools that can boost weak sound levels in a click, automatically ensure that music doesn't drown out speech, and even reduce background noise. We'll show

you how to use those capabilities and give you tips on capturing better quality sound in the first place, such as turning your iPhone into a radio mic or adding a post-production voiceover using a third-party mic. You'll also learn pro split-editing tricks such as adding L and J cuts to help knit your clips together.


How to boost weak sound levels



1 SHOW WAVEFORMS
Click on the timeline's Settings label. In Project Settings, make sure Show Audio Waveforms is ticked. Now you can see the sound levels below each clip as a blue undulating graph. This helps you spot moments of strong or weak sound.



2 ADJUST THE LOUDNESS
Select a clip in the timeline, then click the Audio icon above the viewer. Click Auto to make weaker sound levels become stronger and more audible. Quieter sections of the selected clip's audio waveforms will increase in height.



3 IN THE RED
The peaks of a clip's audio waveform should be yellow. If they peak with a red tip, the sound at that time is too loud. option-click on the waveform's volume bar to add control points. These enable you to dip specific sound levels of sections.


How to perform split sound edits



1 AVOID COMPETING SOUNDS
If you add a cutaway to the timeline, the sound from the layer below may be too competitive. Select the cutaway, then click Volume. Tick 'Lower volume of other clips' to have iMovie automatically dip the sound directly below the cutaway layer.




2 CREATE A J CUT
Control-click on a trimmed clip in the timeline (one with more footage available at its start) and choose Detach Audio. A green bar for the clip's audio will appear. Drag the start of the bar left so that it overlaps the end of the previous clip.



3 MAKE AN L CUT
Repeat for a clip that has unused footage at its end. Extend its audio (or shorten its video) from the right so the sound runs into the next clip. This and the previous step create J and L cuts, and they're used to smooth transitions.

How to get better quality audio



1 BACKGROUND NOISE
To reduce background sounds, click the Noise Reduction and Equaliser button (above viewer). Tick 'Reduce background noise', then use the slider to fine-tune the balance between hearing your subject and losing distracting sounds.



2 ADJUST SOUND QUALITY
The Noise Reduction and Equaliser button also reveals a graphic equaliser. Experiment with the Equaliser pop-up menu's presets to enhance voices, give sharp sounds more bass, and even reduce background hum.



3 ADD A VOICEOVER
By connecting an external mic to your Mac, you can produce pro-quality sound. Cheap acoustic tiles stuck to card absorb room echo for a cleaner sound. Choose Window > Record Voiceover and speak after the countdown.

Sound advice

Create a radio mic.

When filming a subject at a distance from your camcorder, you won't be able to hear them clearly due to the gap between the subject and the camera's built-in mic. Place the mic on your iPhone's headphones near the person's mouth and record a video in the Camera app while they speak. Import the iPhone-sourced clip into iMovie, detach the clip's audio and replace the distant camcorder's inferior sound with the close-proximity sound captured on the iPhone. Now you can hear your subject's voice regardless of their distance from the video camera.






Table of contents in Pages

Adam Banks Automatically add an on-screen or on-page listing to your documents.

A table of contents (TOC) itemises the parts of a document so the reader can jump to the page they want. It could just be a list of chapters, or might include several levels, from sections to individual subheadings.

Pages for Mac can generate a TOC automatically, and the iWork update in March brought this to the iPad. The TOC is based on paragraph styles: apply the same style to each chapter title, heading and so on, and they're collated. You can then choose which to include; the list is limited, so delete unused styles to avoid confusion. If you want

some headings to have a different appearance but take the same role in the TOC, vary their character (not paragraph) styles.

A LA MODE

As usual, matters are complicated by Pages' two modes. In a Word Processing document, having set up your TOC you can insert it into your text to appear on a page. You can then adjust this instance of the TOC, choosing whether it lists everything in the document or only up to the next section break or the next TOC you insert. When you edit the


document, the TOC will update.

In a Page Layout document, the TOC exists as a feature that readers can access on screen; it's included when you export a PDF. You can't insert the TOC into a page for printing, although there is a workaround...


When your Page Layout document is finished, long-tap its thumbnail in Pages' Browse screen and choose Duplicate. Open the copy and convert it to Word Processing mode (see Genius Tip, top right) then create and insert one or more TOCs. Double-tap below the TOC on the page to select it and choose Copy from the contextual menu. Open your original Page Layout doc and paste the TOC into a text box. It appears as plain text formatted with tabs, which you can format. The 'master' TOC will continue to update as you make changes, but those changes won't propagate to the 'print' TOC.

"In a Page Layout document, the TOC exists as a feature that readers can access on screen; it's included when you export a PDF."


How to generate a table of contents




1 IDENTIFY HEADINGS
The table of contents reflects how a document is broken up by headings. This Word Processing document has three levels of subheading (plain, green and blue). The TOC will be based on their paragraph styles.




3 FORMAT APPEARANCE
You can adjust the appearance of each level of heading by selecting one heading, changing its font settings, then tapping the Update button that appears on its paragraph style. The rest change to match.




5 INSERT A TOC
Place the cursor at the start of the document, then tap the paragraph icon (top right of the on-screen keyboard) and choose Page Break. Tap on the new page, tap the TOC icon and choose Insert Table of Contents.



2 APPLY PARAGRAPH STYLES
Apply paragraph styles to headings using the paintbrush icon. Note that you can tap the style name to show a list of styles, then tap Edit to drag styles up and down: make sure they're in correct order of importance.



4 GENERATE TOC
Tap the Table of Contents icon, second from the top left. Tap Edit and choose the paragraph styles of the headings you want to include. Tap Done and you should now see a tabulated list that matches your content.



6 ADJUST YOUR TOC
Tap a line in your TOC, then the paintbrush. You can tweak the font in the Text tab (but can't apply styles). In the Table of Contents tab, set the range of this TOC or tap Customise Styles to change what's included.

Store and sync notes across your PC and devices

Nick Peers examines two solutions to help you record all kinds of notes on your desktop computer, plus access them on all your mobile devices.

Everyone needs a note-taking app – a way of quickly scribbling anything from a to-do list to a detailed breakdown of an upcoming project. When it comes to choosing which app to trust with your notes, you're spoilt for choice. One place you won't find a notes app is in your default Ubuntu installation, however – you'll need to make a choice.

In this tutorial we're going to focus on two note-taking apps, Standard Notes and Joplin. They share some common features: they're both open-source, naturally, both support syncing and work across multiple platforms, so you can access and update your notes from anywhere. But there are also some notable differences, which means you'll

probably find one is a better fit for your needs than the other.

While Standard Notes is open-source it does lock some of its more advanced features behind a subscription-based paywall. However, if you're looking for a simple, text-only note-taking app, there's a lot to recommend it. Your notes are synced and stored online, using end-to-end AES256 encryption to ensure they're kept private.

There are no limits to the number of notes you can store, but that's largely because they're plain text and there's no support in the free version for storing attachments, so your footprint will always be a modest one.

The advantage of storing them online – as you'll see later – is that you can

access them through any trusted web browser as well as your regular desktop and mobile apps. Let's start by installing the desktop version. Head to <https://standardnotes.org> and click the prominent blue download button to save the AppImage to your ~/Downloads folder. Now open a Terminal window to move the image and make it executable with the following commands:

```
# mv ~/Downloads/standard-
notes-3.0.15-x86_64.AppImage
~/standard-notes-
3.0.15-x86_64.AppImage
# cd ~
# chmod a+x standard-notes-
3.0.15-x86_64.AppImage
```


Finally, launch the application – if prompted to integrate the app into your desktop, choose Yes:

```
# ./standard-notes-
3.0.15-x86_64.AppImage
```

The app will launch for the first time from the Terminal. Close this, and then relaunch it from its new shortcut in the application launcher.

You'll be prompted to sign in or register an account – this isn't mandatory if you plan to use the app exclusively on your desktop (in which case click Close). If you do want to use Standard Notes across multiple devices or simply want an encrypted backup automatically stored online, click the Register button to get started.

Enter your email and set a strong password, but note the warning about the lack of a password reset option. This is a security measure, so make sure you've stored it somewhere safe in a password manager like KeePass or BitWarden. Then click Register. When



Getting around Standard Notes

1 VIEWS
Use the left pane to navigate around your notes, applying tags to help quickly setup filtered views.

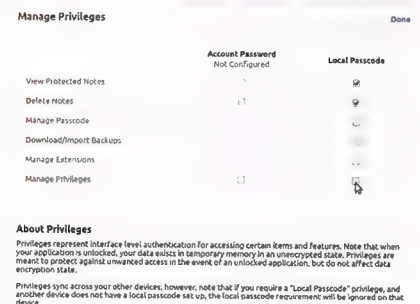
2 ACCOUNT
Click here to access account-related settings, as well as setting a local passcode.

3 VIEW OPTIONS
Click here to change how the selected notes are sorted, plus choose what information to display.

4 NOTE VIEW
The middle pane displays your notes, filtered according to what views you've selected or your search terms.

5 NOTE EDITOR
Click here to edit your notes – the free version of Standard Notes supports plain-text notes only.

6 NOTE OPTIONS
Click the buttons above your note text to access context-sensitive options.



Protect your notes and other Standard Note settings by deploying a local passcode in addition to your account password.

you install Standard Notes on other platforms – including your Android mobile – click Sign In to link it to your account, and all notes will be automatically synced to the new device.

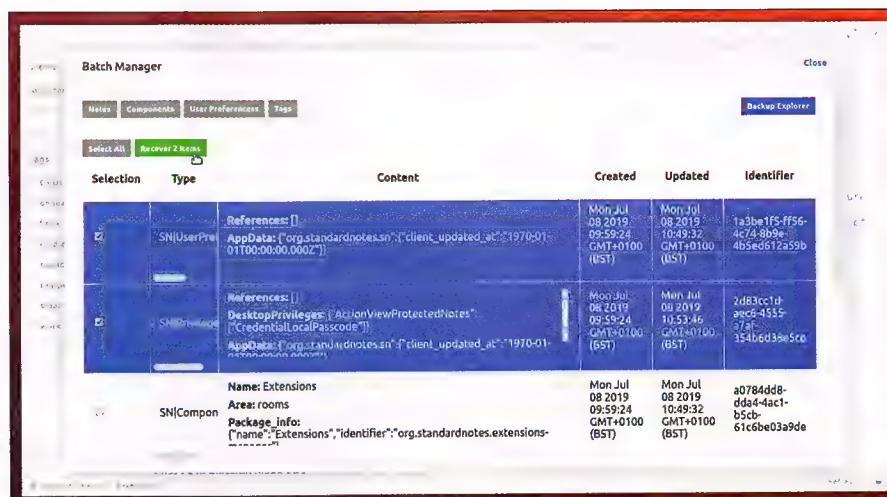
From here it's all pretty straightforward. Standard Notes lets you store multiple plain-text notes, each with its own title. Notes can be assigned one or more keyword tags, which can be used to filter the list of notes shown as well as to provide additional keywords when searching your notes for specific content.

The user interface is a similarly simple affair – the annotation (opposite page) reveals all the key elements that make it up. Pay particular attention to the four buttons in the note editor pane. The Options menu shows options for pinning your note to the top of the list, along with options to archive, lock and protect it. Lock prevents accidental editing (select Options > Unlock to remove it), while Protect allows you to seal the note, requiring either your account password or a locally set passcode to open the note in future.

Before selecting Protect, click the Account button in the bottom-left corner of the Standard Notes window to verify that you're signed in, or to set a local passcode if not. Once done, return to the note and select Options > Protect – the first time you do this, you'll be redirected to the Manage Privileges window where you should tick 'View Protected Notes' and click Done to apply protection to this and all subsequent notes you choose to protect.

Notes are clearly labelled if they're archived, locked or protected, while clicking a protected note displays the passcode or password prompt along with an option to remember your choice for anything from five minutes to a week. You can then easily unprotect the note via Options. Head back to Accounts to reopen the Manage Privileges screen in future to require your password or passcode to perform other sensitive actions, from deleting notes to downloading or importing backups.

The Options button also enables you to set global display settings for the note editor: these include toggling the display font between monospace and proportional, enabling a built-in spell-check and support for adjusting the editor's margins. Next to it is the Editor button, which is where you'd go to switch to a different editor with additional features like support for markdown code and formatting – but this is one of numerous features that require the Extended package, a



You can restore notes and settings stored offline to Standard Notes via its Backup Explorer, part of the Batch Manager extension.

paid-for subscription that's outlined in detail in the box (right). The same is true for the Actions button next to it.

ROLL BACK THE CHANGES

The final button on the note editor – Session History – enables you to view recent changes and roll back to an earlier version of the note if necessary. Click the timestamp entry to preview the older revision, then click Restore or 'Restore as copy' to recreate it as a separate note. By default, only the current session's changes are stored, so once you've closed the app they're lost.

If you select 'Enable saving history to disk', all your changes are permanently stored – but note the warning about the potential effect on performance, which could be a deal-breaker on older machines. Other options under 'Session history' include disabling the auto clean-up function if you want to preserve all changes, however small, plus options for clearing the note history to free up space.

Speaking of backups, Standard Notes automatically syncs all your changes online if you've registered and signed into an account, but if not you should

click the menu button in the top-left corner of the window and choose Backups > Change Backups Location to set an offline folder to store periodic backups of your notes as plain text files.

The contents are unencrypted unless you've set a local passcode, and you can restore backups by clicking Extensions at the bottom of the Standard Notes window, scrolling down and clicking Open under 'Batch Manager'. Click Backup Explorer followed by Choose File to select your last backup, then enter your passcode or password if required. You'll then see a list of all backed-up elements (notes, settings and extensions). Tick the ones you want to restore and click 'Recover x items' to finish. You'll find free apps for other computer platforms as well as Android and iOS mobiles within their respective app stores. Simply log into your account on each device, and your notes will be automatically synced between them. Your data is held securely in the cloud, which also makes it accessible through your web browser on any trusted device. Just browse to <https://app.standardnotes.org> and you'll see an identical user interface to

Extending Standard Notes

In its free form, Standard Notes is a capable text-based notes editor.

But if you're willing to shell out for an Extended subscription, you'll gain so much more. First, you unlock access to almost 30 extensions, including a range of different themes and editors that enable you to introduce word-processor-like formatting to your notes, with support for images, tables and other elements. There's even an editor for creating simple spreadsheets.

The feature list is undoubtedly impressive, and you can road-test all this for free by visiting <https://standardnotes.org/demo>, where you'll get access to a web-based demo account that you can navigate and use as if logged into your own account. The biggest issue is the pricing: currently only a subscription-based model is available, with prices ranging from US\$9.99 a month to US\$149 for five years – which works out at US\$2.48 a month.

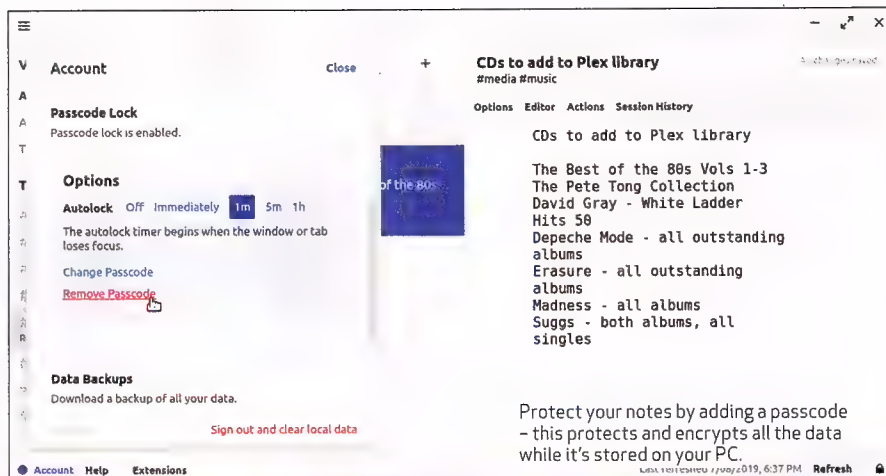
the Linux app, which works in exactly the same way. If you're on a shared device, make sure you don't simply close the browser window when you've finished: instead, take the time to log out via the Account button by clicking 'Sign out and clear local data' to purge the browser's cache.

JOPLIN

Standard Notes has all the core functionality you need for simple text-based notes, but if you'd like to implement formatting like bold and headings, you can avoid taking out an Extended subscription by trying an alternative. Our current favourite is Joplin (<https://joplinapp.org>). It'll appeal to those who'd prefer to retain complete control of the sync platform, but you do lose the ability to access your notes through any web browser – not necessarily a problem, as apps are available for mobile and desktop.

While Joplin doesn't have all the bells and whistles found in the paid-for version of Standard Notes, it does have numerous features not found in that app's free edition, which makes it a worthy alternative. First, its built-in markdown editor enables you to add formatting to your documents, and it also lets you embed images and attach files. You can import notes from Evernote, plus install a web-clipper extension in Firefox or Chrome to capture web pages and screenshots to incorporate into your notes too. There's also a command-line version for using in the shell environment – see the box (below) for details.

As with Standard Notes, you can't



"While Joplin doesn't have all the bells and whistles found in the paid-for version of Standard Notes, it does have numerous features not found in that app's free edition."

install Joplin through the Software Centre. Instead, open a Terminal window and issue the following command:

```
# wget -O - https://raw.githubusercontent.com/laurent22/joplin/master/Joplin_install_and_update.sh | bash
```

This downloads and runs a script that will in turn download and install the

latest version of Joplin. Once done, it'll be accessible from the application launcher, although you may have to reboot before the app icon appears properly.

When you launch Joplin for the first time, it'll display a four-paned window. It's superficially similar to Standard Notes: on the left is a collapsible (click 'Toggle sidebar') pane for choosing your view by selecting a specific notebook or tag. The notebook section reveals a key difference with Standard Notes: Joplin enables you to organise your notes into different notebooks – perfect for separating content by subject.

The selected view's notes are displayed in the next pane, while the final two panes display the markdown editor on the left and a preview window on the right. You'll see that the markdown editor window relies on code to format your text in the right-hand pane, but a series of buttons beneath the note title window simplify the formatting of selected text by inserting the correct code as required. Common keyboard shortcuts such as Ctrl-B for bold are also supported.

Adding further notes and notebooks is simple, using the buttons at the top of the program window, along with a handy 'New to-do' button. This creates a note in the usual way, but displays a tick box next to its title in the notes pane, allowing you to quickly mark it as completed.

There's also a Layout button, this cycles between displaying both editor and preview window side by side or simply displaying each window on its own in turn.

Run Joplin from the shell


If you'd like to run Joplin from the Terminal or other shell environment, you'll need Node.js 8.x or later. If it's not on your system, enter the following commands (skip the first if curl and git are already installed):

```
# sudo apt install curl git
# curl -sL https://deb.nodesource.com/setup_8.x | sudo -E bash -
# sudo apt-get install -y nodejs
Once installed, issue the following two commands:
# NPM_CONFIG_PREFIX=~/.joplin-bin npm install -g joplin
# sudo ln -s ~/.joplin-bin/bin/joplin /usr/bin/Joplin
Once done, type joplin and press Return to launch the application.
```

It runs independently of your desktop app, but has a similar interface, albeit with just three panes: Notebooks, Notes and the currently selected note. Use Tab and Shift-Tab to move between the windows, and the cursor keys to navigate within the currently selected pane.

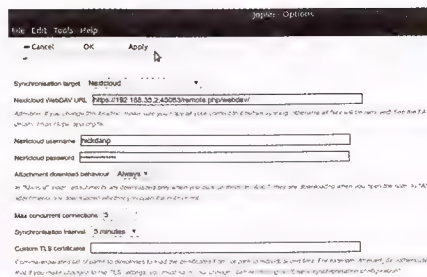
To issue commands, press : followed by the command – for example, type :config editor nano (substitute nano with your choice of text editor) and press Return to associate it with Joplin. Now typing :edit #n enables you to open the note for editing in that editor. Alternatively, press : followed by Tab to cycle through a list of available commands when in command-line mode; refer to <https://joplinapp.org/terminal> for detailed instructions.

Set up encrypted sync in Joplin



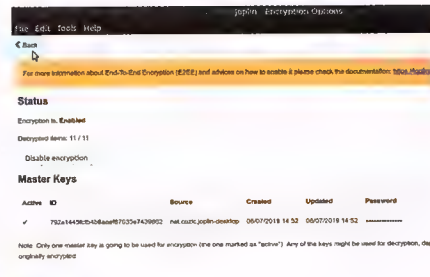
1 CHOOSE SYNC PROVIDER

Select Tools > General Options. First, select your cloud storage from the 'Synchronisation target' drop-down: choose 'File system' and enter its full path if you have a shared network folder ready. OneDrive and Dropbox options are easy to set up: choose the option, click OK and then click Synchronise in the main window, before following the instructions to connect Joplin to your account.



2 NEXTCLOUD PRE-REQUISITES

If you select Nextcloud, enter its WebDAV URL (visit your Nextcloud account in your browser and click the Settings button in the bottom left corner to find this), followed by your Nextcloud account username and password. If your Nextcloud account is on a secure (https) connection, tick 'Ignore TLS certificate errors' before clicking Apply followed by 'Check synchronisation'.



APPLY ENCRYPTION

Once sync has been tested successfully, choose 'Tools > Encryption options' and click 'Enable encryption'. Choose a strong password to protect your notes and click OK. You'll see your master password listed: verify it's marked as okay. Click Back followed by Cancel to return to the main Joplin screen. Sync runs every five minutes, or you can click Synchronise to perform a quick manual sync.

GET IN SYNC

By default, Joplin stores everything locally on your PC. It supports multiple devices and sync, but these need to be applied manually using your own network or cloud storage as an intermediary. Five choices are available: Dropbox, OneDrive, WebDav, Nextcloud or a network share. The step-by-step guide (opposite page) shows how to set this up, plus establish encryption to ensure your notes are kept private – an essential step if you're syncing to a publicly accessible cloud service. Note that WebDav and Nextcloud sync is unreliable between mobile and desktop at present, so you're better off using OneDrive or Dropbox if you plan to sync with an Android or iOS device.

Once you've set up and performed your first sync, Joplin will continue to synchronise in the background at five-minute intervals. In the meantime, it's time to switch to your new device. Configure synchronisation in exactly the same way as you did with your PC, pointing to the same provider and account via the configuration screen. If you're an Android user, tap to bring up the menu to select Synchronise and then log in to your account if necessary.

Joplin will bring in your synced items as a separate notebook, to sit alongside whatever items are locally stored on your new device. You'll see they're all listed as 'encrypted' and unreadable. To resolve this, simply tap 'Press to set the decryption password', enter the password and tap Save. You'll also need to encrypt any existing notebooks on the device before they sync into the cloud. Set up a master password for

these, which you'll need to enter on your PC before you can access those notebooks on there.

The Joplin web-clipper makes it easy to grab content from websites to incorporate into your notes. Get started by opening Joplin and selecting 'Tools > Web clipper options'. Click 'Enable Web Clipper Service' and then click the relevant link for Firefox or Chrome to download and install the add-on.

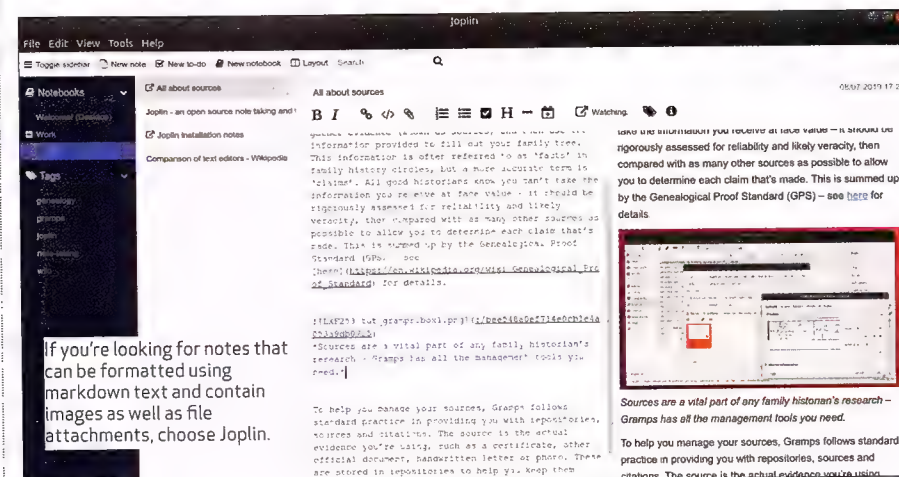
Once in place – which might take a while, as we found it claimed not to be working for a few minutes before magically starting – make sure Joplin is running in the background. Then simply browse to a website and click the Joplin icon in your browser toolbar to bring up a pop-up menu of options.


First, choose what you want to clip: a simplified version of the page, a complete HTML version of the page, the selected portion of the page only, a screenshot image of the page, or the

URL. Make your choice and a preview will appear. Next, choose which notebook to save the note to, add any required tags (as you type, it will offer to autocomplete from existing tags in Joplin) and edit the title in the Preview box before clicking Confirm.

Once created – you may have to wait a few seconds while the page is downloaded and any necessary conversion is performed – switch back to Joplin where you should find your note waiting for you.

Both Standard Notes and Joplin are great note-taking tools, with different sets of features to appeal to different audiences. Those looking for a simple, fuss-free solution with minimal effort will most likely lean towards Standard Notes, but if you find yourself yearning for its premium-only features on a budget, you'll find Joplin fills the gap admirably – without costing a cent. ■





Dive right in

Choose a development environment



Get started with Augmented Reality on Android

Augmented Reality is fast becoming more affordable on mobile devices. Darren Yates looks at the AR world according to Google, plus the software and hardware you'll need.

Like it or not, *Pokemon Go* ended up becoming one of the most downloaded gaming events of 2016 and the first mainstream app to bring the world of augmented reality (AR) to our collective mobile consciousness. Sure, AR tech has been with us since at least the 1960s and you could argue since WWII, thanks to the gyro-powered gunsight reflector used in the famed Spitfire fighter plane. But until *Pokemon Go*, AR was mostly reflected in standalone hardware – think Microsoft HoloLens or Google Glass. But over the last 12 months or so, we've started seeing new AR capabilities reach right down into surprisingly affordable Android devices. Throw in new free software development kits from Google and the world of AR looks set to go far beyond chasing pocket monsters.

HOW MOBILE AR WORKS


In broad brushstrokes, AR on your phone typically works by taking the view from your rear camera and

augmenting it with a 3D view of some virtualised 3D object. The trick is that the 3D view maintains perspective in alignment with your real-world view using a technique called 'motion tracking'. This works so that if you rotate your phone around the 3D object, the 3D perspective view of that object follows your real-world view.

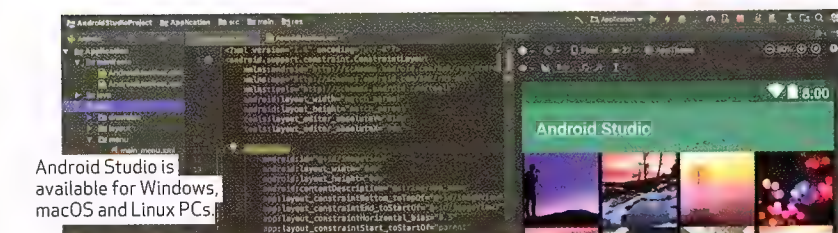

But if you lift the bonnet and consider all the major components in AR from a processing perspective, there's an awful lot going on – the CPU has to process live video from the rear image sensor, calculate the 3D perspective view of your virtual object and overlay that view onto the initial video image in the right place and all in real time. Then it has to maintain the visual lock between the two views via motion tracking – this involves reading the phone's accelerometer, gyroscope and magnetometer sensors to understand the spatial positioning and ensure the 3D object is in the right place and at the right size. Smartphones are computers, but this is pushing the friendship just a tad.



Samsung's Galaxy A30 gets you ARCore functionality for under \$400.



Android Studio is the default IDE for coding your own Android apps.



Android Studio is available for Windows, macOS and Linux PCs.

WHAT YOU NEED

Nevertheless, while it does push smartphones pretty hard, AR is increasingly reaching down into the 'mortgage-friendly' phone market, so if you're interested, here's a short list what you need:

- ARCore-ready Android phone
- A half-decent PC (quad-core recommended)
- Google ARCore/Sceneform for Android software development kit (SDK)
- Google Play Services for AR app (Google Play, automatic install)

ARCORE-READY ANDROID PHONE

Like most innovations, this new-found 'DIY AR' capability comes on the back of newly-released higher-performing smartphones. However, what used to be limited to models strictly from the top drawer is fast sliding down the food chain into the sub-\$400 mid-range market. Australia doesn't get access to every ARCore-ready phone, but Google keeps a global list of models capable of performing the 'mathematics' fast enough to make AR possible. The list at

developers.google.com/ar/discover/supported-devices at time of writing covers a decent list of current and recent-past phones, including the Nokia 6.1/7.1/8.1 range, the LG G-series (starting at the G6), plus Samsung's Galaxy A-series, beginning with the 2017 Galaxy A3.

As for 'ARCore', this is Google's new software development kit (SDK) platform, a series of programming instructions that handle much of the low-level work of creating the augmented view, such as reading the sensors, detecting real-world surfaces and more.

The key thing to remember, though, is that not every phone is ARCore-ready and you must consult the Google support list before you buy an AR phone. In most cases, ARCore requires a minimum of Android 8.0, although anything newer would be great. Also, the more CPU horsepower you can throw at it, the better all-around. As a rough rule-of-thumb, ARCore-ready phones appear to need at least eight Cortex A53 CPU cores. The Cortex A53 is ARM's economy-performance CPU

core and you'll generally find four of them in most entry-level pre-paid phones, with eight-core CPUs finding their way into 2018-era phones up to about the \$350 mark.

From a quick look, the cheapest 2019-era ARCore-ready phone we've seen at time of writing is Samsung's \$379-priced Galaxy A30, with its Exynos 7904 chip featuring two Cortex A73 'power' cores and six Cortex A53 'economy' cores.

A HALF-DECENT PC

While ARCore certainly simplifies AR code development, it still is code development, so it won't be for everyone. Still, don't let that stop you. You will, however, need a reasonable PC – although, that depends on how you run your AR code. The ideal way is to run your code on an actual Android phone. This is how I do it and I'm using an ageing 2012-era Intel Core i5-2300 PC running 16GB of RAM and Windows 8.1, with code running on a Samsung Galaxy A30 phone. If, however, you choose to run your code on an 'Android Virtual Device' (AVD) virtual machine instead, you'll need a faster CPU, more storage and plenty of RAM. Personally, I'd suggest at least 16GB of RAM is needed in any serious development situation, but virtualisation is a known RAM hog, so it's important to have plenty and help keep your PC's solid-state drive (SSD) in shape. While we're on it, an SSD with plenty of free space would be ideal.


The key software you'll need is Android Studio (developer.android.com/studio). This is the default Android integrated development environment (IDE) that lets you write your Android code, compile it and load it on your phone for testing. It's 700MB-plus to download before you even begin adding any extra SDKs. Android Studio now supports two programming languages – the original Java, plus the more recent Kotlin. There are similarities between the two, but you can stick with Java if you prefer. You'll find versions of Android Studio for Windows, macOS and Linux-based systems.

GOOGLE ARCORE SDK


The secret sauce that makes this all work is in two parts – the first is the Google ARCore SDK you download and use within Android Studio. The second is the 'Google Play Services for AR' app you install on your AR-ready phone from Google Play (play.google.com/store/apps/details?id=com.google.ar.core). This second part was originally also called 'ARCore' (just to be confusing), with the recent name-change to fix the ambiguity. You can choose to install this yourself, but Google has said this AR-focused services app will be automatically



Sceneform's 'augmented image' example adds 3D frame corners to an image.



Sceneform's chroma-key example uses a green-screen MP4-video animation.



The 'hello sceneform' example lets you place a 3D Android icon on a surface.

installed onto eligible devices as required.

What makes ARCore more useful is the simple ready-to-go example projects you can load and compile, install and run on your ARCore-ready phone. Everyone has to start somewhere, so even if you're not a confident programmer, you can still learn to load up the examples, run them, follow through the code and learn how the functions work.

The examples we're looking at are included in the 'Sceneform for Android' SDK, a Google ARCore 3D framework

that combines the core AR tasks with all of the difficult 3D perspective work, making AR app development much easier. You'll find all of the ARCore SDK downloads at <https://developers.google.com/ar/develop/downloads> and the Sceneform SDK on GitHub at <https://github.com/google-ar/sceneform-android-sdk>.



WHAT TO DO

To get started, you need to install Android Studio onto your system – at least version 3.1 is required, but the new 3.5 release would be ideal. Once installed, download and unzip the ARCore/Sceneform for Android SDK onto your PC and ensure the latest 'Google Play Services for AR' app is on your AR-ready phone. Next, launch Android Studio, choose 'File/Open' from the main menu, navigate to your unzipped 'sceneform-android-sdk' folder, select the 'samples' subfolder and choose the 'hellosceneform' project.

This is Google's AR answer to 'hello world', the first app every programmer

builds in a new language, except this is rather better than the usual printing of 'hello world' to the screen. Plug your phone into your PC via USB, press the green 'play' triangle button top-right, then choose your Android device from the pop-up 'Select Deployment Target' list. This begins the compilation process and installs the code as an app (.apk) on your Android phone. When the app starts, it automatically launches your device's rear camera and you'll see a hand icon making circles on the screen. This is your cue to point your camera towards a flat surface and move your phone in slow circle movements. This allows your phone to detect the surface and the perspective view – it'll show this as a series of white spots on the detected surface. Once these appear, move your phone around and those spots should maintain perspective. When you're ready, tap the screen and you'll get a mini 3D Android icon appear on that surface.

Now this is where the fun begins – physically move around the 3D icon and



it'll change perspective as you move. Not only that, you can tap-and-hold the icon and drag it, as well as add new icons by just tapping a blank spot on the perspective surface.

Sure, this example may not do anything practical, but by drilling down into the code in Android Studio, you can begin to nut out how the code works, modify it and develop your own AR apps.


3D MODELING

Many of the AR examples you'll find (not just in Sceneform but elsewhere) are built around 3D models. At this stage, you create these models in third-party software and import them into Android Studio. The 3D asset formats supported include the common .obj and .fbx files, which you can create with various 3D modelling apps. These include Blender and 3DS MAX, but Google forums indicate you should be able to use any digital content creation (DCC) toolset that creates compliant files. The caveat is Google has been ironing out the bugs and you may find the odd model doesn't import correctly, which is why you should keep updated with new releases of Sceneform. To make it easier, there's an ARCore SDK release for Unity, the popular real-time 3D development platform (<http://unity.com>) available on the ARCore download page above.

CHROMA-KEY VIDEO

While 3D model-based AR is popular, Sceneform/ARCore also supports the use of video, including chroma-keyed video. Chroma-keying is where an object is shot on a uniformly-coloured background that does not appear in the object itself (think 'green screen'). Later, the background colour is replaced by a separate video stream via chroma-keying in software. Think of almost any superhero movie and most scenes will involve chroma-keying.

Sceneform includes a 'chromakeyvideo' example, which is an extension of the 'hellosceneform' example. It replaces the icon model with short video of a stylised cartoon lion



drawn on a green-screen background. When it appears on your phone screen in AR, the green-screen is replaced by the view from the rear camera. Like the previous example, you tap the screen to display the AR view.

AUGMENTED IMAGES

However, tapping the screen isn't only one way to trigger an AR experience. You can also use an image as a trigger. The Sceneform 'augmentedimage' example does this. Load the example, open the project folder in Android Studio's left navigation panel and click on /app/assets/default.jpg – an image of planet Earth. Compile the app, install it into your phone, scan the Earth image on your PC screen with your phone camera and once recognised, the app will place picture frame corners around the image. Move your phone around and the frame corners will track the corners of the image.

In this example, Sceneform has created an image database in your app and the app continuously scans the initial rear camera view looking for any of the images in the database. Once one of those images is detected, the AR function is activated as coded.

AR MARKETING

If you're struggling to consider how this might be used (aside from Pokemon Go), you only need to look to your local supermarket. The issue of giving away plastic toys aside, Australian

supermarket chain, Coles, has been doing a roaring trade with its 'Little Shop' promotion of miniaturised products. So much so, that it's created 'The Coles Little Shop 2 App' on Google Play, a proprietary AR app that shows a different dancer for each of products that make up the promotion. The app requires an Android 6.0 or later phone where you scan the miniature product, follow the prompts and view an animated character busting out some moves. Not your thing? It's not exactly mine, either – but this is an example of new 'augmented reality marketing' and for better or worse, you can just about guarantee seeing more of it over the next year or so.

GIVE IT A GO

There's no doubt AR can be one of the most impressive things you'll see on a phone, but the fact that easier-to-code AR is now possible on a \$379 phone shows just how far smartphones have come in the last decade. AR app development isn't quite at the point-and-click stage just yet and you'll likely always need good coding skills. However, the fact that Google has put together just about all of the tools you'll need to master basic AR for free is impressive in itself.

If you want to get ahead of the curve and know what's coming up in tech, Google Scholar and Google ARCore are good places to start. ■

Create a Raspberry Pi OLED monitor

Les Pounder shows us how to monitor a Raspberry Pi 4 with a small screen and a little bit of Python code.

You need

- Any Raspberry Pi
- An i2c OLED screen (128 x 64 pixels)
- 4x female-to-female jumper wires
- The Code: <http://bit.ly/lxf254i2c>

This issue's project is a tiny OLED screen that will provide useful data about a Raspberry Pi at a glance. We will also run our project on boot to monitor the system as we work.

The OLED screen connects to the Raspberry Pi GPIO using an I2C interface. This is a special interface that only needs four pins on the Pi in order to send data. To connect your Pi please refer to the diagram on the following page.

With the physical connection made we now need to enable I2C on our Pi – do this in the Raspberry Pi Configuration tool. We will also need to install two Python libraries for this project. The first library is for our OLED screen and the second enables our code to work with network interfaces. In a terminal enter the following command.

```
# sudo pip3 install Adafruit-SSD1306 netifaces
```

WRITING THE CODE

To write the code we shall be using Thonny, the default Python editor on the latest Raspbian. You're free to use any editor you want, but please save the code as `live-data.py` – and save often. We start the code with a line indicating where our Python 3 interpreter is located. Our project will be a background app in the system, so this enables it to run.

```
#!/usr/bin/env python3
```


We now import three libraries to add features to our code. The first enables us to use the OLED screen with Python. The second is used for accessing system information in Python. The last is used to access network interface information. Note that we import 'netifaces' but refer to it as 'ni' for ease of use.

```
import Adafruit_SSD1306
import psutil
import netifaces as ni
```

Three more imports, and they are all from the Python Imaging Library (PIL). These will handle drawing content on the screen.

```
from PIL import Image,
ImageDraw, ImageFont
```

In order to use the OLED screen, we need to tell Python where it is, and what size it is. For this we create an



This tiny screen is easy to work with and really cheap. It can provide all the data we need at a glance, without firing up a terminal.

object called `disp` and in there we use the Adafruit library, specifically the `128_64` function. Note that the screen does not have a reset pin (`rst`) so we set that to `None`. Then we begin the connection to the screen.

```
disp = Adafruit_SSD1306.SSD1306_128_64(rst=None)
disp.begin()
```

To configure the screen we need to know the width and height; luckily, the function handles that. These are saved to variables, and then the screen is set to 1-bit colour (black and white).

```
width = disp.width
height = disp.height
image = Image.new('1',
(width, height))
```

Preparing the screen for output, we create an object to which we can draw.

```
draw = ImageDraw.Draw(image)
```

Now we set the confines of the screen. The `padding` variable enables us to stay inside the screen; we set the top to match the padding, and the bottom of the screen as the height minus the padding. We also create a variable, `x`, to keep track of where text will appear on the horizontal axis.

```
padding = -2
top = padding
bottom = height-padding
x = 0
```

Here we gather the data for our screen. The first data collected is the

CPU's temperature, and for this we open a file in read-only mode. This file stores the data from the CPU temp sensor, which we save to a variable called `f`.

```
f = open("/sys/class/thermal/thermal_zone0/temp", "r")
t = f.readline()
```



As the data is stored as a string, in order to do any maths with it we need to convert it to a `float`. The value stored is far too large, so we divide it by 1,000 and round the returned value to one decimal place, giving us the correct temperature.

```
t = float(t)
t = round(t/1000,1)
```

To get the CPU usage percentage, we create a variable called `cpu` and in it store the output from the `psutil.cpu_percent()` function. Then we round the value down, with no decimal place for ease of reference.

```
cpu = psutil.cpu_percent(interval=1)
cpu = round(cpu)
```

To get the RAM usage we again use `psutil`, specifically the virtual memory function. This will return a list of data, and from the list we need the second value, which is available memory. To get the second item we use `[1]`, as a list starts from 0. Then we do a little maths on the value in memory, which represents bytes. We need to convert it to megabytes so we divide the



Only four wires are needed for this project. It's easier to wire than an LED!

Running the code on boot

Our project is best used on boot, so here are the extra steps to make that happen. First we need to make our code executable, and for that we need to open a terminal and type

```
$ chmod +x live-data.py
```

Now to make the code available to everyone as a command, we need to copy it to /usr/bin, so in the terminal type

```
$ sudo cp live-data.py /usr/bin/live-data
```

Now give the command a try. In the terminal type the following and watch the OLED screen.

```
$ live-data
```

So now we have the command, let's use it on boot. In the terminal we shall use crontab and edit cron to start the command on boot, and run the command every minute that the Pi is on.

```
$ sudo crontab -e
```

In crontab, go to the bottom and add this line.

```
* * * * * live-data >/dev/null 2>&1
```

Now save by pressing Ctrl-X and pressing Y, then Return. That's it! Now you can reboot your Pi and watch your OLED screen as every minute your Pi is on, it will give you live data. The spring evenings will just fly by...

value by 1,024, twice.

```
memory = psutil.virtual_
memory()
memory = memory[1]
memory = round(memory / 1024
/ 1024)
```

The last piece of data is the internal IP address of the Ethernet port. For this we create a variable called `ip` and in it we store the data obtained from `netifaces`

`ifaddresses` function. This requires the interface name as an argument, and in the returned data we need to go into a series of nested lists: first at position 2, then inside that list, position 0. Then we use the `addr` key to retrieve the data stored in a dictionary.

```
ip = ni.ifaddresses('eth0')
[2][0]['addr']
```

The next step is to add an alert for our CPU. If the temp (`t`) is over 50C, it will change the text on the screen. Otherwise we get `NOMINAL`.

```
if t > 50:
```

```
    ALERT = "CPU TEMP ALERT"
```

```
else:
```

```
    ALERT = "NOMINAL"
```

We now start writing text to the screen. The first two lines write the CPU temp and usage.

The first line is our CPU temp and you will see that we start writing in the top left (`x`), two pixels from the top of the screen (top is -2). We write a string of text, and convert it to a string `t`. Then we use the default font, and set the colour to 255 (white). The second line is the same, but we set the move down eight pixels.

```
draw.text((x, top),
"CPU Temp is "+str(t)+"C",
font=ImageFont.load_default(),
fill=255)
draw.text((x, top+8), "CPU
usage is "+str(cpu)+"%",
font=ImageFont.load_default(),
fill=255)
```

The next three lines are our Alert status, available RAM, and IP address.

```
draw.text((x, top+16),
ALERT, font=ImageFont.load_
default(), fill=255)
draw.text((x, top+25),
str(memory)+"MB Free",
font=ImageFont.load_default(),
fill=255)
```

```
draw.text((x, top+33),
"eth0:"+ip, font=ImageFont.
load_default(), fill=255)
```

The very final two lines of this code will update the display and show the live data.

```
disp.image(image)
disp.display()
```

Save the code and click Run; your OLED screen will come to life and show the latest data for your machine. See below to have this code run on boot. ■

How and where you use machine-learning

Understanding machine-learning is great, but where do you use it? Darren Yates looks at some practical examples that make use of this maths-meets-computing platform.

We're all pretty familiar with the concept of autonomous vehicles, whether you're a fan of driverless cars or the Boston Dynamics 'Atlas' robots that leap over obstacles with human-like agility (bostondynamics.com/atlas). But unless you have your own car corporation or financing your own squad of robots, it can be difficult to get a handle on just how machine-learning is applicable in your own applications or your own business. This month, we're looking at some practical applications of machine-learning and how it fits into the bigger world of knowledge discovery.

CASE STUDIES

One of the best ways to learn how to apply knowledge in a particular area is to read how it's been used previously. Now you might get the impression that a lot of academic research is just theory, but that's actually not the case – a huge amount of it is also applying that research into practical applications. What's more, much of that practical research is available online. Google Scholar (scholar.google.com.au) is Google's search site for research, including peer-reviewed published papers and new unpublished 'pre-print' work. While a lot of the published work is 'pay-per-view', there's plenty that is freely accessible via download, going back to Alan Turing's ground-breaking paper on 'Computing Machinery and Intelligence' and beyond.

AGRITECH

One of the areas where Australia is leading in practical research is agriculture technology, more commonly known as 'agritech'. According to the Australian Bureau of Statistics, Australian agriculture output topped \$59 billion in 2017/18. Yet, with drought and demand for greater yields, farmers are increasingly turning to technology to provide answers. Unsurprisingly, this has made the agritech business one of the fastest growing areas, not just in Australia, but around the world. The amount of work involving machine-learning in agriculture might surprise you – everything from satellite imagery

The screenshot shows the Agerris website with a red tractor in a field. Text on the website reads: "REVOLUTIONISING AGRICULTURE AND TRANSFORMING ON-FARM OPERATIONS WITH ROBOTICS, AI AND INTELLIGENT SYSTEMS". Below this, a caption says: "Start-ups like Agerris are using ML and robotics to boost." Below the website screenshot is a Notepad window titled "wifi_localization (1).txt" containing a table of RSS readings.

-64	-56	-61	-66	-71	-82	-81	1
-68	-57	-61	-65	-71	-85	-85	1
-63	-60	-60	-67	-76	-85	-84	1
-61	-60	-68	-62	-77	-90	-80	1
-63	-65	-60	-63	-77	-81	-87	1
-64	-55	-63	-66	-76	-88	-83	1
-65	-61	-65	-67	-69	-87	-84	1
-61	-63	-58	-66	-74	-87	-82	1
-65	-60	-59	-63	-76	-86	-82	1
-62	-60	-66	-68	-80	-86	-91	1
-67	-61	-62	-67	-77	-83	-91	1
		-61	-67	-72	-86	-81	1
		-61	-65	-73	-84	-84	1

First seven columns are RSS readings, the last column the room number.

to predict crop yields through to automated robotics. In Australia, new start-ups like Agerris (agerris.com) are developing intelligent robotics to help with boosting crop efficiency and livestock welfare, reducing costs and improving sustainability.



FOLLOW THE DATA

Understanding how machine-learning works in practice can be a bit daunting – if this is your first exposure to what goes on 'under the bonnet', it's probably all just a sea of data and computers and algorithms and stuff. As we alluded to last month, much of machine-learning is about turning that sea of data into knowledge you can do something useful with. So, a great place to start is to look at good data available and see how it's created. The University of California, Irvine (UCI) has a machine-learning dataset repository well-thumbed by

researchers, containing datasets on a huge array of topics from medical research to car evaluation. The key is to look at how the data is arranged, as it can tell you plenty about the kind of information researchers are looking for – it should help give you an idea of how to round up and use data you need in your application.

MEDICAL RESEARCH

Many medical diseases are being researched with the help of machine-learning and the UCI data repository gives you a glimpse at how researchers are using data and what they're looking for. One of the earliest datasets is the 'breast cancer wisconsin' set ([archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic))). This 1990-era dataset from the University of Wisconsin contains 699 records of breast tissue samples, each with nine



"If this is your first exposure to what goes on under the bonnet, it's probably all just a sea of data and computers and algorithms and stuff."

attributes recording various aspects of the tissue cells, followed by a classifying or 'class' attribute with diagnosis values of either 'benign' or 'malignant'. The goal of machine-learning used here (called 'pattern recognition') was to try and find an accurate method for matching the cell attributes with the cancer diagnosis to enable faster and more accurate detection. Since then, there wouldn't be a field in medical research that hasn't been enhanced by machine-learning, whether its disease detection, drug efficacy or patient care.

WI-FI LOCATION DETECTION VIA SMARTPHONE

Finally, here's an example a little closer to home – UCI has a dataset called 'wireless indoor localisation' (archive.ics.uci.edu/ml/datasets/Wireless+Indoor+Localization). It contains 2000 records, each with eight attributes, seven of which contain the received signal strength (RSS) values from seven Wi-Fi signals observable on a smartphone. The last attribute is the 'class' or classifying attribute that details which of four rooms (1-4) the

phone was in at the time of those seven readings.

So what's the problem this data is trying to solve? Phone location detection is pretty easy if you're outdoors – you can either use GPS or triangulate signal strength with known phone network towers. Getting location indoors is a lot harder. This dataset accompanies a research paper presented in 2017 that attempts to use machine learning to identify room location by RSS. With 16,000 data points (2000 x 8), it's impossible to find the rules that determine the relationship by hand – that's where machine-learning comes in (the research showed that using a particular form of neural network, the researchers could predict the room location with better than 95% accuracy).

EXAMPLES OF CLASSIFICATION


In both of these examples, the goal is to identify the patterns that relate a series of attributes of a record (spreadsheet row) to a particular category or classifying value (column). This common machine-learning task is called 'classification' if the category values are discrete (yes/no, days of the week etc) like our examples. If, however, they're numeric (such as salary, population etc), the task is known as a 'regression' analysis.

THE RIGHT QUESTIONS GET THE RIGHT DATA

Machine-learning has a bucketload of hype surrounding it at the moment, almost to the point that it can do anything. The thing is, though, it's a bit like a hammer. Hammers are great for hammering in nails, not so good for screws. In order for machine-learning to work its magic, you have to feed it the right raw material – and that's good quality data. Every business has data – that's no problem. The problems will be deciding what do you want to find out, do you have the right data, if not, how to get the right data and finally, when you've got the right data, is it in the right form you can use.

A key to obtaining good-quality data is almost always 'domain expertise' – employees who work at the coal face are the ones who'll often provide you with the ground-truth that allows you to categorise events or results. It's also why asking questions to elicit responses that match up with the data is so important – and often, such hard work.

Machine-learning is a powerful tool that can be applied to any range of applications – but usually the real trick is understanding how, where and when to use it. ■



Want to know the future of machine-learning? Start with Google Scholar.

downtime

WHAT'S HAPPENING IN THE WORLD OF FUN.



Lego Star Wars Battles is a new mobile strategy game that draws from the entire saga

MIX AND MATCH VEHICLES AND CHARACTERS FROM ACROSS THE FRANCHISE.

Warner Bros. Interactive, The Lego Group, and Lucasfilm announced that *Lego Star Wars Battles* will arrive sometime in 2020.

"We hope players will enjoy the unique combination of *Lego Star Wars* with deep strategy in deck building and innovative territory capture by constructing Lego towers," said TT Games Brighton studio head Jason Avent.

Xbox is getting new Sport Blue and Night Ops camo controllers and you can pre-order them now

THEY'RE UP FOR PRE-ORDER ON AMAZON.

Looking for the perfect new controller to fill out your collection of Xbox One gamepads? What about two new controllers? Microsoft recently unveiled a duo of new official, special edition designs for the Xbox Wireless controller coming this fall, and you can put your pre-order(s) in for them on Amazon right now. There's the Xbox Wireless Controller – Sport Blue Special Edition and the Xbox Wireless Controller – Night Ops Camo Special Edition.

Deadly Premonition 2 is the sequel we didn't dare hope for and it's coming to Switch in 2020

Francis York Morgan returns in a not-quite-*prequel* adventure.

Deadly Premonition 2: A Blessing in Disguise is a video game that exists somehow, thank goodness, and it's coming to Nintendo Switch. The sequel to the cult classic B-game will continue the story of detective Francis York Morgan next year as he takes on a new case in 2020 – actually an old case, more on that in a bit – but don't worry if you haven't played through the original yet. *Deadly Premonition* is coming to Switch too, and you can start playing it today as *Deadly Premonitions: Origin* on the Nintendo Switch eShop.


If you're never heard of *Deadly Premonition* before, you have been missing out. It first came out in 2010, and its intrinsic weirdness that strode the line between *Twin Peaks* inspired weirdness, janky video game weirdness, and warmhearted comedy weirdness instantly captivated fans of, well, weird games. If you're not sure if it sounds right for you, just watch this cutscene about the "Sinner's Sandwich" and you'll know.

Cyberpunk 2077 is getting some "multiplayer action" post launch, CD Projekt Red confirms

MULTIPLAYER MODE IS OFFICIALLY COMING TO CYBERPUNK 2077.

CD Projekt Red has officially confirmed multiplayer is in the works for *Cyberpunk 2077*. The studio took to Twitter to reveal the news that some "multiplayer action" will be added some time after the initial release on 16 April, 2020. Just like *The Witcher 3: The Hunt*, *Cyberpunk 2077* will also be getting some free DLC and additional single-player content post-launch, but we don't know what this will include just yet.

It's long been suspected that *Cyberpunk* might get some kind of multiplayer modes somewhere down the line. CD Projekt Red has been pretty hush hush about it all so far, and has continued to maintain its focus on talking about the single-player campaign. Just a month or so ago, CD Projekt Red still appeared to be weighing its options about including multiplayer. The studio's senior concept artist Marthe Jonkers, who spoke to *Eurogamer*, said there were no guarantees it would be included.



Submarines are sailing into *World of Warships*

David Hollingworth rigs for periscope depth and goes hands-on with a whole new way to battle.

World of Warships dev Wargaming is no stranger to tweaking and adding content to its popular military titles. Submarines, a wholly new, fourth class of vessel, have been announced, and the game will never be the same again. Wargaming's already invested years into the new addition, with the first hint of this direction coming with last year's Halloween event, where players took command of fantastic steampunk submersibles in a limited, strictly PvE event. The work that went into even that short test was immense, and the work's not stopped since.

Submarines are going to be a whole new threat. Submarines operate at three preset depths: on the surface, at periscope depth, and fully submerged, and must balance an ever-dwindling oxygen resource that depletes ever faster the deeper they get. As always, getting the feel of the submarines right is a balancing act between accuracy, balance, and player expectation. The deeper a submarine is, the harder it is to spot, and at depth only destroyers have a chance of even


coming close to finding them – which is another whole new system being added to the game. Historically, destroyers were often tasked as sub-hunters, armed with groundbreaking sonar technologies and a brutal array of depth charge ordnance.

This is modelled in the game by what the devs call a new 'mini-game'. As destroyers are tooling around the map, large circles will appear on the water's surface, denoting that there is a sub somewhere in that volume; the destroyer player will have a short amount of time to enter the circle, after which a second circle will appear not too far away, which the player must then get to again. This repeats three or four times, with each time the volume getting smaller and smaller, as the destroyer's sonar effectively narrows down where the enemy sub is. When the destroyer reaches the final circle, it automatically launches depth charges.

Submarines are, of course, not without their own powerful attack. From periscope depth they can launch torpedoes much like any other ship –

though only fore and sometimes aft – but once your fish are in the water sub drivers can launch a special, non-damaging attack that Wargaming is calling an 'acoustic wave'. It's a little hand-wavey at the moment, and not entirely historically accurate; each acoustic wave hit (and you can hit targets both fore and aft for added effect) makes your torpedo attack more likely to hit, as the torps home in on the target. And let's just say subs are not the toughest boats. Even a destroyer's guns are a direct threat; a battleship's broadside, if it finds you before you can dive, is going to give you a very short, sharp illustration in explosive dynamics.

This whole new aspect of the game is still a work in progress – and the trick is to integrate submarines into the game in a way that doesn't so much change the way *World of Warships* plays, but to add a new and entirely organic element – one that feels like it belongs there. Full integration of sub-surface warfare is expected in 2020. ■



\$56.95 | PC, XB1, PS4 | REMNANTGAME.COM/EN

Remnant: From the Ashes

A surprisingly nourishing mechanical soup.

This odd genre mash-up looks like it was coughed out by a game development algorithm. It's equal parts a third-person looter-shooter and a *Dark Souls*-inspired boss rush, with a cooperative emphasis and a dynamically generated campaign. Stylistically it's inspired by fantasy, science-fiction, Westerns, post-apocalyptic worlds and various kinds of horror fiction. Even the title reads like it was generated by a Twitter bot.

And it gets off to a rather underwhelming start. Its introduction sees your character washed up on an island with nothing but a sword and the clothes on their back. Eventually you end up at Ward 13, the Firelink Shrine of *Remnant*'s world. Here you learn that a race of tree-people called the Root have conquered Earth. Despite the fact that you've just arrived, the leader of the

Ward decides you're the person to stop them.

While *Remnant* resembles a cover shooter, in practice it plays much more like *Dark Souls* from an over-the-shoulder perspective. Basic enemies are aggressive and hit hard, so to stay alive for any length of time you need to evade their attacks. Gradually, you pick your way through the ruins, moving between glowing crystals analogous to *Dark Souls*' bonfires and fighting creatively designed bosses.

Unlike *The Division*, which drip-feeds the player incrementally better guns, *Remnant*'s loot is less common and more tailored. The real treasures are dropped by bosses, powerful items which can be taken back to Ward 13 to craft new weapons and weapon mods.

Remnant's most unique feature is its dynamically generated campaign.

Environments can be laid out in different ways and dungeons located in different places. Similarly, each boss encounter has several possible variants. I played through the first half of *Remnant* twice, and in the second run, Singe was replaced by the Ent, a giant tree with a face like Cthulhu. It's a clever way of encouraging you to invest in other players' campaigns, and arguably better design than having a massive, shared world you only explore once.

Elsewhere, although *Remnant*'s dimension-hopping makes for a fun adventure, the game struggles to tie its threads together narratively. *Remnant* borrows *Dark Souls*' approach to storytelling, using item descriptions and fragmented conversations with NPCs to build its lore. But *Dark Souls*' world is cogent, consistent, and


painstakingly built. *Remnant* too often feels like it's making things up as it goes along. There are plot reasons for this revealed right at the end, but it's a bit of a cop-out, veering dangerously close to one of the most wearisome clichés in storytelling.

Nonetheless, I'd happily play more of *Remnant*. Its strange worlds, robust challenge, and bespoke loot make it a refreshing alternative to the drab militarism of *The Division*. ■ Rick Lane

Verdict

A flawed but fascinating hybrid, this is a looter-shooter that's charged with potential.





\$21.50 | PC, LINUX, MAC | DICEYDUNGEONS.COM

Dicey Dungeons

One RNG to rule them all.

A lazy journalist would call this *Hearthstone* meets Yahtzee, and just as a stopped clock tells the right time twice a day, sometimes lazy journalist comparisons are bang on. *Dicey Dungeons* rolls you a few dice, presents you with cards which harness those dice in damage-dealing or character-buffing fashion, then pits you against a series of opponents in a dungeon-crawling RPG setting. It feels like playing a CCG while wearing X-ray specs.

Because while the roguelikes and card games that *Dicey Dungeons* is referencing with its visuals and mechanics would hide away their RNG elements, Terry Cavanagh's new game puts it under a spotlight. Quite literally. The setup is that of a gameshow, so while you're fighting sneezy hedgehogs and sharply dressed hammerhead sharks there's a studio audience gasping and cheering over it all. They ooh and aah when you open chests in its simple dungeon maps. There's even a wheel of fortune.

Within the opening hour you're dealing out shock,



fire, and freezing damage, locking opponents' dice, and firing off double-attacks when your limit break bar fills up from taking damage. And crucially, you feel like some kind of mastermind tactician while you do it. That's the trick in *Dicey Dungeons*. You're always being given the chance to make a smart decision, whether it's holding off on doing damage to pool some dice into a healing card or using buffs to hit exact numbers required by status effect cards.


There's a deceptively massive chunk of game here, and if the likes of *Hearthstone* or *Gwent* never clicked with you it's an easy way to slide into the world of agonising over deck builds. You only have space for a

handful of cards here but striking the right balance is just as engrossing.

Like an over-eager Tinder date, the vibe is very much casual and lighthearted at first but before too long you're made aware that the expectation is absolutely that of a long-term prospect. Unlike a stranger you swiped right on, *Dicey Dungeons* doesn't keep liking your Instagram posts months after you stop playing. ■ Phil Iwaniuk

Verdict

Endearing, compulsive, and just weird enough to keep you wanting to explore its world.



Man of Medan

Makes clichéd horror feel fresh again.

\$49.95 | PC, XBI, PS4 | BANDAINAMCOENT.COM

If it were a movie, this would be the kind that releases direct-to-video, one of those schlocky DVDs lining the convenience store displays. As a film, it would be fun, but familiar and forgettable. But as co-op adventure game, *Man of Medan* is unparalleled. Think of it like a movie where you take control of one character at a time, oscillating between quicktime action scenes and dialogue sequences. Occasionally you'll get time to look for story clues, tools, and soak in the setting.

Early into my online co-op playthrough, my coworker Joanna and I coerce our ship captain into letting us explore a sunken plane at the bottom of the ocean, but we don't play two characters in the same scene. We play different scenes simultaneously.

Half of *Man of Medan* keeps Joanna and I separated, playing unique interconnected scenes. Based on the clues we discover and the decisions we make, we're able to both avoid disaster and steer headlong into it, depending on what we choose to share. Failure is usually fun, a necessary part of horror. But death isn't always ceremonial, often cut short from one failed quicktime event or a misread button prompt. ■ James Davenport





Enter The Matrix

Jeremy Peel recalls this flawed attempt to reinvent movie games.

Twenty years ago a new science fiction world was born, as relatable as it was outlandish. *The Matrix* was a rallying cry for anybody whose internet life had become more real to them than their daytime drudge; a freeing fiction that overlaid our own world like a sheet of graph paper and made sense of its unexplained repetitions and irregularities. As much as it was a nightmare or warning, it was also vicarious fantasy – it was pasty Keanu Reeves, asleep at his PC while Massive Attack plays in his headphones, clinging to the side of a building for dear life, eventually bending reality to his will. On some level, everybody watching wanted that call from Morpheus. To wake up, to know kung fu, to move like they do.

That's what *Enter The Matrix* offered. Its title was a mission statement. Rather than mimic scenes from the movies, you

could become a part of their story, interwoven with the green strands of its code.

It was born from a time when Hollywood was crowning a new wave of powerful directors, a younger generation to whom the artistry of videogames was self-evident. Peter Jackson recognised a fellow auteur in Michel Ancel, and tasked the *Beyond Good & Evil* designer with adapting *King Kong*. Guillermo Del Toro was losing himself in *Ico*, which he declared a masterpiece. And the Wachowskis, suddenly at the head of a multimedia empire, used their position to work with as many of their favourite artists as they could. They commissioned Neil Gaiman to write a comic; in Tokyo, Madhouse and Studio 4°C produced anime. For the interactive counterpart, they went to the developer behind a goofy sci-fi shooter – MDK.



HOLLYWOOD BOUND

Shiny Entertainment's games had long been distinguished by their imagination – from the squishy, gross-out cast of *Earthworm Jim*, to the hybrid RTS strangeness of *Sacrifice*. A strict no-sequels policy ensured oddball ideas could thrive, making the studio a spiritual precursor to *Double Fine*. But since the beginning, founder David Perry had his sights set on Hollywood – hence the 'Entertainment' suffix, which he believed would make the company appear more worldly to film production companies.

He got his wish, and then some, with The Wachowskis – who came to game development with a story already written, and heads full of ideas for specific sequences, down to the level design. It sounds stifling for a studio used to flights of fancy, and there's no doubt whose vision



This'll teach you not to stop and search so much.



Jada Pinkett Smith earns her top billing by oozing cool.



Neo might survive a fall like this. Niobe won't.

Enter *The Matrix* delivers on – as a longtime Shiny Entertainment fan in 2003, I played through the entire game without once clocking the developer behind it.

That tight directorship, however, is the very reason it's a game worth returning to. *Enter The Matrix* formed part of a bold experiment to connect *The Matrix*'s many spin-offs to the movies.


This was not like the expanded universe of *Star Wars* – the mass of licensed novels, cartoons, and games so sprawling and disconnected that Disney felt comfortable dismissing the lot from canon in 2012. Instead, the Wachowskis oversaw an interdependent story that lent each part equal importance.

SWITCH IN PERSPECTIVE

For *The Animatrix*, the Square studio behind *Final Fantasy: The Spirits Within* created the *Final Flight of the Osiris* – a photorealistic CG short that followed a doomed hovership crew much like Neo's. Their discovery, of a huge machine army drilling down into Zion, precipitated the events of *The*

NOT THE ONE

You think anyone gets a cool hacker name?



Matrix Reloaded.

Then, Shiny's game ran parallel to the film, embellishing and expanding on plot points referenced on-screen. And after the conclusion of the trilogy, a Monolith-developed MMO officially took the story forward into the future – at least until 2009, when *The Matrix Online* was shut down.



The Wachowskis wrote and directed an hour's worth of cutscenes for *Enter The Matrix*, and it's those that still represent

its key draw. For once a game adaptation was not a poor relation to its cinematic inspiration. Instead it was merely a switch in perspective, from the Nebuchadnezzar to another crew caught up within the machinations of Morpheus.

Or at least, that was the idea. The very first level of *Enter The Matrix* reveals a tension in that premise. You're planted in a US post office to pick up a tape containing the last transmission of the Osiris, so that you can hand it over to Zion – literally passing the baton from one medium to another. Yet, before you can leave that post office, the shutters come down and you're pushed into a lobby firefight, pinwheeling out from behind pillars as bullets strip the marble from the walls.

It's a clear echo of the standout action scene from the original film (for which Perry turned down a game deal, not anticipating the phenomenon that would follow).

Enter The Matrix was a product of two different urges – one to dive from the rooftop after the established spectacle of the



movies, and another to follow its own path. For the most part, it gets away with attempting both, presenting different angles on sequences familiar from the movies.

When Niobe catches Morpheus on the bonnet of her car during the climactic freeway fight of *The Matrix Reloaded*, the moment functions nicely as a surprise twist. Niobe has volunteered her crew to go chasing after Morpheus rather than stay behind to defend Zion, so being there to save his life validates her choice.

Playing the game afterwards, however, fleshes out that arc with powerful effect. As Niobe or

Ghost, you drive that car up the freeway, fending off cops and Agents so that you can show up for Morpheus – knowing exactly how important your role will be.

It's an approach that gels nicely with the Matrix sequels' less Messianic take on The One. If Neo is fated to save mankind, the story goes, it's not a foregone conclusion – it'll only happen with the help and sacrifice of others. As in the Wachowskis' transmedia empire, every part is of equal importance.

TIGHT SCHEDULE

Oddly, there are other times in this shared story where it feels as if cinema-goers have been

robbed of spectacle so that *Enter The Matrix* players can enjoy it instead. At one point in *Reloaded*, Morpheus sets up a heist of sorts.

Niobe's crew is tasked with blowing up a power plant to prevent a security system being triggered – but the movie cuts straight to the explosion. It's only *Enter The Matrix* players who wind their way through the plant, shutting off the lights and taking down snipers in the dark. It might be the one time in history that a film has been compromised for the benefit of its licensed game, rather than the other way around.

The game commits to this split-perspective story still

JACKING IN Making sense of the complicated Matrix chronology



Bits and Pieces of Information
Catalogues the conflict between humans and machines in comic form.



The Matrix
In which The One turns from scruffy office drone to spoon-bending superhero.



Final Flight of the Osiris
CG short which sets up the sequels' existential threat to mankind.



The Matrix Reloaded
Neo battles with rogue programs and his own purpose.



The Matrix Revolutions
The machines reach Zion and battle ensues. As does utter confusion.



The Matrix Online
This MMO officially continued the series, but it's dead now.

Anybody still connected to the Matrix is a potential Agent

further by offering an opening choice between Ghost and Niobe as protagonist. While Ghost has a nice sideline in Hume quotes, Niobe's love triangle with Morpheus and Zion's commanding officer makes her the more intriguing choice – not to mention her red faux-crocodile skin trench coat, by some distance the most detailed and mesmerising thing in the game. The pair's shared path diverges at certain points, usually when collaborating on something. As a teenager playing Ghost, I distinctly remember shooting out the front wheel of an aeroplane during takeoff with a sniper rifle. Now, as Niobe, I find myself sprinting through hangars to catch up with the grounded vehicle.

It's these high level, structural quirks that are still worth celebrating in *Enter The Matrix*. The fighting's decent enough, too. You've got a finite amount of bullet time which is topped up by kills – chain enough together and you'll have a constant trickle of slo-mo throughout the fight. Ammo is scarce, pushing you to swap often between

shotguns, uzis, and a basic take on martial arts. In short, it keeps you moving, just as Neo and Trinity did in that lobby.

What's more, Shiny smartly committed to contextual moves long before that became the industry standard for action games. A sidestep towards a wall while in bullet time becomes a cartwheel, during which you can empty half your clip or switch to another weapon. At its best, it's smooth and instinctive.

That said, it can be tempting to forgive old games too much in lieu of putting them in their proper context. You know what else came out in 2003? *Max Payne 2*. Remedy had outmanoeuvred the Wachowskis, making their aesthetic interactive years before it was officially licensed. Shiny, meanwhile, never quite left the '90s – too often designing levels as roomy, blocky interiors rather than the spatial combat puzzles the third-person action genre was perfecting.


The developer clearly suffered, too, from the hard limits of their project. While the Wachowskis threw money and support at *Enter*

The Matrix, it still had to come out when *The Matrix Reloaded* did, imposing a tight two-year development schedule that left the game janky. It was a depressingly familiar problem for a project that had otherwise broken free from the low expectations of licensed movie games.

Afterwards, Shiny worked on *The Matrix: Path of Neo*. Though built on the same combat, it was the opposite of *Enter The Matrix* – a greatest hits compilation that pulled all of the most famous scenes from the films. The Wachowskis had capitulated to fans who simply wanted to waggle a controller while the movies played out once more before their eyes.

By contrast, *Enter The Matrix* stands as a flawed monument to ambition, and perhaps even a model for developers and directors to follow and improve upon in the golden age of the cinematic universe. ■





RAD


Neon flair can't make up for painfully average action.

\$29.95 | PC, XBI, PS4, SWITCH | BANDAINAMCOENT.COM/GAMES/RAD

Double Fine's post-apocalyptic mutant-basher embraces the current 80s retro fad. The devs have poured lavish detail into every synthesised riff. Every sunset in the background is a neon CRT wonder, every bark of the narrator a gruff, biting commentary. Vendors will lavish you in free gear for landing a sick 720. It's the *Goonies* meets *Bill & Ted* in an end-of-the-world adventure.

It's the far future. We've wrapped up one apocalypse, but another hit us with a cheap shot. With the adults nowhere to be seen and mutants breaking down the walls. Rather than picking up new tricks and tools from vendors or drops, you're mutating a new power every time you collect enough rads. But you're getting properly messed up by these mutations. Each comes with a horrifying animation of your kid writhing in pain, as some fresh hell grows out of your helpless form.

Alas, I've been dealt many a killing blow from foes that drop fire upon death, because I couldn't drag my poor kid out in time. Frankly, hitting things feels rubbish in *Rad*. It's a killer flaw, even as the game tries to shine. ■ Natalie Clayton



US\$29.99 | PC, XBI, PS4, MAC | REBEL-GALAXY.COM

Rebel Galaxy Outlaw

Makes up in style what it lacks in depth.

Despite the vast potential of space's expanse, coming from a small dev *Rebel Galaxy Outlaw* knows its limits, restricting its galaxy to 40-odd solar systems named after cities in the US. Its guitar-twanging rock soundtrack and vibrant colours give it the tone of a punky western. It's thrifty where it needs to be, and does a commendable job of embellishing its relatively rote gameplay loop with presentational swagger.

Driving the game's spunky spirit is the protagonist Juno, a tough vagabond on the trail of the people who killed her husband and tried to knock

her off too. You gain most of your intel by chatting with the patrons of bars dotted around the various space stations of the cosmos.

Things get tough out there in the coruscating cosmos. It may look like a neon-tinged music video set in the *Guardians Of The Galaxy* universe, but it can be unforgiving, and I find myself dying often in my ungainly cargo container of a starting ship. Combat sticks within the norm of dogfighting space sims, with a few well-advised convenience controls and snippets of flair raising it a little bit beyond the mean. You have the usual

assortment of blippy lasers and blasty missiles, but one of the better features is auto-follow, activated by holding the left trigger. Not only is this great for orientation, but turn up your thrusters and you can really get up close to enemy ships, chasing them down as you dart between clusters of space debris and freighters.

There are hints that you can play *Rebel Galaxy Outlaw* as a plucky trading vessel, with a dynamic fluctuating market that lets you trade goods, as well as the higher-risk activity of ferrying counterfeit goods to those willing to buy them. But realistically, you'll spend the bulk of your spacefaring in combat. This can get a little repetitive, but with the carrot of ship upgrades and better weapons dangling over you, it offers enough incentive to keep pushing on.


■ Robert Zak



Verdict

A lively space-western that sacrifices other games' scale in favour of more style.





\$18.50 | PC, X81, PS4, SWITCH | HEADSPUN.CO

Headspun

Manage the offices of a broken brain.

Headspun is another twist on the genre, casting you as the boss of Cortex, a company headquartered inside the mind of Theo, a man emerging from a five-week coma after a car accident. Theo's memories are scattered all over his subconscious, most of Cortex's employees are dead, and it's up to you, CEO Ted, to get his brain back in order. This manifests as a management game in which you rebuild the company, interposed with live-action video of Theo's convalescence in hospital.

Ted wakes up from his own coma to find Cortex in ruins following the accident, albeit kept in some semblance of order by his business partner Teddy, who represents Theo's creative, romantic side. The game progresses in 24-hour days, with time passing whenever you engage Theo in mini-games.

Your interactions with the live-action part of the game are through Cortex Control, the main room in Cortex where a screen shows you what's happening in Theo's hospital room each day through his eyes. Beyond the neurologist, the only significant character you



interact with is Jack, Theo's estranged best friend who believes that the two of them have a chance of publishing their own comic series.

It's a reasonably engaging story, as you cross-reference the things you're told in the live-action snippets with the information you piece together in Cortex's Memory Bank. Information surfaces about Theo's journey from aspiring comic-book artist to wanker banker, as well as details that make you doubt the veracity of Jack and Teddy's words.

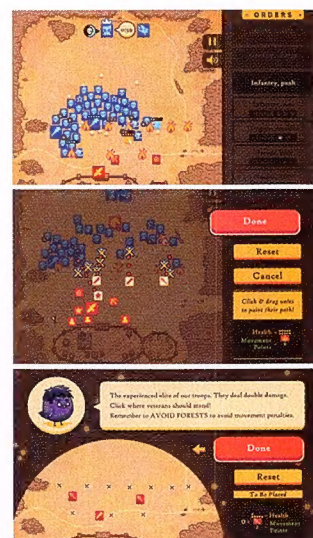
From Cortex Control, you also assign workers to sidequests, which usually involve rebuilding rooms in Cortex, researching new activities for Theo, or side-quests which clean up various parts of the

building. The live-action snippets are few and far between, meaning that you spend most of the game playing mini-games that build up Neuro-Credits, which you use to hire staff and rebuild rooms. The mini-games – basic maths, crosswords, timing challenges, and other forgettables – are repetitive and mundane, though, and make this game thoroughly un-compelling.

■ Robert Zak

Verdict

A high-minded effort and solid story undermined by bugs and braindead mechanics.



Winterfalling

Fight the big battle.

FREE | BROWSER-BASED | RARYKOS.ITCH.IO

Game of Thrones' Battle of Winterfell was a thrilling fantasy bust-up that had armchair strategists pulling their hair out, over the apparently dumb tactics involved. Now it's your turn to do better. There are no Dothraki or giants here, as every unit type is represented by cute little icons on a tactical map, but it's a cheekily similar situation to the battle present in the show. You're attempting to preserve an important figure against a relentless hoard of undead enemies, and you do this by placing units at strategic positions, and by yelling out a few limited orders after the battle starts. But your most important decisions happen before fighting starts. Do you hide non-warriors inside the crypt, to improve morale, or make use of them during the battle? Where do you place your archers, your cavalry, your hero units? "So you hated the TV battle," *Winterfalling* asks, "how exactly would you have done it differently?" It turns out that surviving Winterfell is a lot harder than it looked in the show.

■ Tom Sykes





3,000 person game of Hide and Seek shut down by police


Unofficial Facebook event snuffed out by shop owners.

Social media platforms seem to have finally found their true calling – bringing the world together for fun, ridiculous events. Last month it was the Storm Area 51 Facebook festival, this month it's a 3,000 person game of Hide and seek at an Ikea in Glasgow, Scotland. Perhaps because of its less-deadly nature, the latter actually made it to fruition in September, with Scottish hipsters and internet trolls turning up to tuck themselves in some Scandy furniture. The store caught wind of the unsanctioned fun-time and Police Scotland posted five officers to rebuff any youths looking too young to be spending their Saturday shopping for furniture.

Tesla drivers becoming known for falling asleep at the wheel

HOPEFULLY NOT A HIGHWAY TO HELL.

Despite pretty clear legal signalling by the company and more than a few fatal slip ups, many Tesla owners still seem determined to put their lives in the hands of their self-driving cars. The latest incident in the ongoing saga is a video posted to Twitter that shows a Tesla 'driver' and their passenger asleep at the wheel on a freeway in Boston Massachusetts. The Model X that was driving is supposed to have safeguards to ensure drivers have their hands on the wheel 'at all times' so it can disengage if it isn't capable of managing the situation, but the driver in the video has seemingly circumvented, or is able to simply sleep through, those alarms.



Google maps satellite image shows location of murder victim

1997 COLD CASE BODY FOUND USING A GOOGLE MAP PICTURE.

A resident of Wellington, Florida noticed an unusual looking object in a Google Maps satellite image of the lake at the back of his house. The Grand Isles gated-community resident reported to the authorities that there was a car submerged in the lake after inspecting the site further with a drone. Upon retrieving the vehicle the Palm Beach County Sheriff's department discovered a skeleton inside it, which was later identified as William Moldt, a missing persons case from 1997. At the time the missing persons report was filed the housing estate was still under construction.




Russian gun-wielding robot is being sent to the ISS

THIS HUMANOID ROBOT IS ALSO WELL VERSED IN PHILOSOPHY.

Despite probably not needing the services of a Russian gun-slinging robot that can argue the meaning of life, the ISS inhabitants just had a brief visit from such a humanoid machine. Russian deputy Prime Minister Dimitri Rogozin insisted earlier in the year that Skybot F-850, known pejoratively as Fedor, was "not a Terminator" after video footage of it shooting handguns at targets made headlines around the world. Fedor's space activities were a little less spicy than its land based hobbies (it helped the cosmonauts with chores), but it did do the important job of test flying a new shuttle and booster combination for the Russian space agency Roscosmos. ■

SMARTER COOKIES ENABLED



JOIN US ONLINE AT
WWW.PCGAMER.COM

NBN SMB Solution for Ethernet WAN (FTTP, HFC, Fixed Wireless, Satellite)

DrayTek
Aust & NZ
www.draytek.com.au



50x
VPN

Vigor2926 Series

Dual-WAN broadband router with Gigabit LAN, Security Firewall, 802.11ac (AC2000) Wi-Fi, VoIP and 50 x VPN tunnels including 25 x SSL-VPN tunnels

- Dual Gigabit Ethernet WAN ports for Failover and Load-Balancing and High Availability mode
- Two USB 2.0 ports for connection to two 3G/4G LTE USB mobiles, FTP server and network printer
- 4 x Gigabit LAN ports with multiple subnets and 50,000 NAT sessions
- 50 x VPN tunnels (including 25 x SSL-VPN tunnels) with comprehensive secure protocols
- Fast VPN throughput, VPN load-balancing and backup for site-to-site applications

- 16 x VLANs for secure and efficient workgroup management
- IPv6 & IPv4
- IEEE 802.11ac (AC2000) wireless Access Point; dual band; up to 1.7Gbps throughput (Vac/ ac model)
- 4G LTE SIM card slot (L model)
- VoIP with 2 x FXS and 1 x FXO Line Port (V model)
- Object-based SPI Firewall and CSM (Content Security Management) for network security
- Supports VigorACS 2 Central Management System for remote management
- Central VPN Management for 8 remote Vigor routers
- Central AP Management for deployment of multiple wireless VigorAPs
- Free Smart Monitor Network Traffic Analyzer for 30-nodes

Vigor2952 Series

Dual-WAN broadband Fibre Router with Security Firewall, 4 x PoE LAN ports (power budget up to 60 Watts), 100 x VPN tunnels including 50 x SSL-VPN tunnels



TAIWAN
EXCELLENCE
2016

100x
VPN

- 1 x Gigabit Ethernet/SFP combo WAN port, 1 x Gigabit Ethernet port
- Multi-WAN Load Balancing & Failover and High Availability mode
- 4 x Gigabit PoE Ethernet LAN ports with 60,000 NAT sessions (P model)
- Complies with 802.3at Power over Ethernet Plus (P model)
- Supports PoE power up to 30 watts for each port, and total PoE power budget up to 60 watts (P model)
- 1 x USB port (USB3.0) support 3G/4G LTE USB mobile
- 100 x VPN tunnels including 50 x SSL-VPN tunnels with Central VPN Management
- VPN Load Balance and Redundancy
- Object-oriented Firewall with Content Security Management and QoS
- IPv6 & IPv4
- Multi subnet WAN/LAN through 802.1Q
- 1 x USB port (USB2.0) support file sharing, external storage, network printer or thermometer
- Central AP Management; Central Switch Management; Smart Monitor Traffic Analyzer
- Supports VigorACS 2 Central Management System for remote management

Vigor2133 Series

Broadband router with Gigabit LAN, SPI Firewall, 802.11ac (AC1200) Wi-Fi, VoIP and 2 x VPN tunnels including 2 x SSL-VPN tunnels



2x
VPN

- 1 x Gigabit Ethernet WAN port
- 2 x USB ports for 3G/4G access, FTP server and Network printers
- Load Balance and Failover functions for WAN and USB 3G/4G WAN
- 4 x Gigabit Ethernet LAN ports with 30,000 NAT sessions
- IPv6 & IPv4
- NAT throughput over 450Mbps
- Object-based SPI Firewall, Content Security Management (CSM) and QoS
- 2 x VPN tunnels including 2 x SSL-VPN tunnels
- 802.11ac (AC1200) Wi-Fi, up to 300 + 867Mbps speed
- SIP VoIP (2 x FXS) for Vac model
- Central AP Management for deployment of up to 2 wireless VigorAPs
- Supports VigorACS 2 Central Management System for remote management
- 2 years back to base warranty



i-LAN Technology Pty Ltd trading as DrayTek Aust & NZ

www.i-lan.com.au
www.draytek.com.au

(02) 9838 8899

sales@i-lan.com.au

© Copyright 2019 i-LAN Technology Pty Ltd. All rights reserved.